Journal of Computer Security 0 (0) 1
IOS Press

arXiv:1603.00087v1 [cs.CR] 29 Feb 2016

Effective Sequential Protocol Composition in Maude-NPA

Sonia Santiago [a,c], Santiago Escobar [a], Catherine Meadows [b], José Meseguer [c]
[a] *DSIC-ELP, Universitat Politècnica de València, Camino de Vera, s/n 46022 Valencia Spain*
E-mail: {ssantiago,sescobar}@dsic.upv.es
[b] *Naval Research Laboratory, 4555 Overlook Ave SW, Washington, DC 20375, USA*
E-mail: meadows@itd.nrl.navy.mil
[c] *Department of Computer Science, University of Illinois at Urbana-Champaign, The Thomas M. Siebel Center for Computer Science 201 N. Goodwin Ave. Urbana, IL 61801-2302, USA*
E-mail: meseguer@cs.uiuc.edu

Abstract. Protocols do not work alone, but together, one protocol relying on another to provide needed services. Many of the problems in cryptographic protocols arise when such composition is done incorrectly or is not well understood. In this paper we discuss an extension to the Maude-NPA syntax and its operational semantics to support dynamic sequential composition of protocols, so that protocols can be specified separately and composed when desired. This allows one to reason about many different compositions with minimal changes to the specification, as well as improving, in terms of both performance and ease of specification, on an earlier composition extension we presented in [18]. We show how compositions can be defined and executed symbolically in Maude-NPA using the compositional syntax and semantics. We also provide an experimental analysis of the performance of Maude-NPA using the compositional syntax and semantics, and compare it to the performance of a syntax and semantics for composition developed in earlier research. Finally, in the conclusion we give some lessons learned about the best ways of extending narrowing-based state reachability tools, as well as comparison with related work and future plans.

Keywords: Cryptographic Protocols, Formal Verification of Secure Systems, Sequential Protocol Composition, Protocol Verification, Maude-NPA

1. Introduction

The area of formal analysis of cryptographic protocols has been an active one since the mid 1980's. The idea is to verify protocols that use cryptography to guarantee security against an attacker —commonly called the *Dolev-Yao* attacker [13]— who has complete control of the network, and can intercept, alter, and redirect traffic, create new traffic on his/her own, perform all operations available to legitimate participants, and may have access to some subset of the longterm keys of legitimate principals. Whatever approach is taken, the use of formal methods has had a long history, not only for providing formal proofs of security, but also for uncovering bugs and security flaws that in some cases had remained unknown long after the original protocol's publication.

A number of approaches have been taken to the formal verification of cryptographic protocols. One of the most popular is model checking, in which the interaction of the protocol with the attacker is

0926-227X/0-1900/$27.50 © 0 – IOS Press and the authors. All rights reserved

symbolically executed. Indeed, model-checking of secrecy (and later, authentication) in protocols in the bounded-session model (where a *session* is a single execution of a process representing an honest principal) has been shown to be decidable [33], and a number of bounded-session model checkers exist. Moreover, a number of unbounded model checkers either make use of abstraction to enforce decidability, or allow for the possibility of non-termination.

It is well known that many problems in the security of cryptographic protocols arise when the protocols are composed. This is true whether the composition is parallel, in which two different protocols are executed in an interleaved fashion, or sequential, in which one or more child protocols use information from executing a parent protocol. Protocols that work correctly in one environment may fail when they are composed with new protocols in new environments, either because the properties they guarantee are not quite appropriate for the new environment, or because the composition itself is mishandled. Security of parallel composition can generally be achieved by avoiding ambiguity about which protocol a message belongs to (as in, e.g. [26,9]). The necessary conditions for security of sequential composition are harder to pin down, since they depend on the guarantees offered and needed by the particular protocols being analyzed.

To see an example of the problems that can arise, consider the analyses of the Internet Engineering Task Force's (IETF) Group Domain of Interpretation (GDOI) protocol [3], in which the third author of this paper was involved. GDOI is a secure multicast protocol built on top of the IKE Version 1 (V1) [27] key distribution protocol, which had already undergone at least one formal analysis [28] and substantial peer review by the IETF. However, early versions of GDOI had two subtle flaws, one arising from the composition of GDOI with IKE, and the other arising from the way a subprotocol of GDOI was used by the parent protocol. One was a type confusion attack that took advantage of the fact that IKE V1 headers began with random numbers instead of a field indicating what type of protocol it was [30]. An intruder could take advantage of this confusion to obtain a group key to which it was not entitled. The discovery of this attack led to a redesign of the GDOI protocol before it was submitted as a standard. Another attack involved a subprotocol of GDOI, called the *Proof of Possession* (POP) protocol. The GDOI specification was not clear about the situations in which POP was to be used. Once these were clarified, it was discovered the POP was also subject to an attack [29]. This discovery led to a modification of GDOI to fix this vulnerability.

The importance of understanding sequential composition has long been acknowledged, and there are a number of logical systems that support such compositional reasoning. The Protocol Composition Logic (PCL) begun with [15] is probably the first protocol logic to approach composition in a systematic way. Logics such as the Protocol Derivation Logic (PDL) [6], and tools such as the Protocol Derivation Assistant (PDA) [2] and the Cryptographic Protocol Shape Analyzer (CPSA) [12] also support reasoning about composition. All of these are logical systems and tools that support reasoning about the properties guaranteed by the protocols. One uses the logic to determine whether the properties guaranteed by the protocols are adequate. This is a natural way to approach sequential composition, since one can use these tools to determine whether the properties guaranteed by one protocol are adequate for the needs of another protocol that relies upon it. Thus, PCL and the authentication tests methodology underlying CPSA are used to analyze key exchange standards and electronic commerce protocols in [10] and in [23] respectively, via sequential composition out of simpler components.

Less attention has been given to handling sequential composition when model checking protocols, especially in the case in which an instance of a parent protocol can spawn multiple instances of subprotocols, e.g., in the case of a master key agreement protocol that can be used multiple times to generate a session key. We believe that this is an imbalance that needs to be corrected, for logical systems and

state exploration techniques make complementary contributions to our understanding of the security of a protocol. Logical methods allow us to construct proofs from basic assumptions, much as we develop protocols that use basic cryptographic algorithms. These logical systems provide insight into how a protocol achieves security, and what basic assumptions it depends on. State exploration tools, on the other hand, provide concrete attacks that can be used in fixing a protocol. Moreover, they are also useful for discovering behaviors that, while they may not violate specified security properties, nevertheless turn out to be undesirable. This can be used to inform and refine the logical systems. Finally, state-exploration-based models can provide a useful semantics for logical systems.

The problem is in providing a specification and verification environment that supports composition. This is not necessarily straightforward; we note in particular that the two leading formal calculi underlying most current cryptographic protocol analysis tools, the pi calculus [1], and strand spaces [20] do not provide general sequential operators that can be used to specify protocol compositions. That does not mean of course that one cannot define in these languages protocols that are actual compositions of other protocols; it just means that sequential protocol composition is not supported at the language level and must be encoded by the user in ways that may depend on the particular composition at hand. The aim of this work is to provide specification primitives for a wide variety of compositions at the language level in a way that is both transparent to the user and sound and complete with respect to a desired semantics.

There are several ways that composition can be handled in state exploration systems. One is to not modify the tool at all, but to handle everything at the specification level, by concatenating protocols that are being composed: for example, a master key agreement protocol followed by a session key distribution protocol. This requires no modification, but besides being tedious to specify when many different ways of composition are possible, it also cannot be used to represent cases in which a parent protocol can have an arbitrary number of children, as is indeed the case in the master/session key case. Another is to compose protocols at execution time, but without modifying the operational semantics of the tool, an approach we took in [18]. Although this minimizes the modifications made to the tool, it can lead to counterintuitive and inefficient methods of specification and analysis, since we are adapting composition to the original semantics instead of the other way around. Finally, we can extend the operational semantics of the tool, but minimize such an extension by adding or modifying as few semantic rules as possible.

In this paper we describe how this third approach has been applied to the Maude-NPA protocol analysis tool and its strand-space-based semantics. We first give an abstract composition semantics, first introduced in [18], that extends Maude-NPA's operational semantics using the concept of *parameterized* strands [25] augmented with a separate composition operator. We then describe an extension of Maude-NPA to Maude-NPA *with composition* via *synchronization messages*, in which composition is achieved by means of strand space parameters, but without the separate composition operator, and prove soundness and completeness of the operational semantics of Maude-NPA with composition with respect to a subset of the abstract semantics. This extension allows us to minimize the changes made to Maude-NPA, as well as the number of rewrite rules that need to be added to its operational semantics. We provide evidence that this approach to extending Maude-NPA is probably optimal by comparing it with an earlier approach we took in [18], in which composition was implemented via *protocol transformation*, in which the synchronization messages implementing composition were passed along the Dolev-Yao channel. Although this required fewer modifications to Maude-NPA, since it already supported communication in the Dolev-Yao model, the additional communication overhead had a negative impact on performance. We illustrate this via experiments comparing the performance of both composition approaches, via synchronization messages and protocol transformation.

Thus the contributions of this paper are the following:

1. It provides a formal definition of sequential protocol composition in the strand space model (Section 5).
2. It provides a new operational semantics of protocol composition in Maude-NPA (Section 5).
3. It provides a simple and intuitive syntax for protocol composition in Maude-NPA (Section 5).
4. It describes an implementation of protocol composition directly in Maude-NPA via the operational semantics, giving a proof of its soundness and completeness with respect to a subset of the abstract semantics. (Section 6).
5. It provides an experimental evaluation of the new operational semantics, and compares its performance with respect to the protocol transformation technique presented in [18] (Section 8).

The rest of the paper is organized as follows. In Section 2 we introduce two motivating examples of sequential protocol composition, which will be used throughout this paper as running examples. After some preliminaries in Section 3, we give an overview of the Maude-NPA tool and its operational semantics in Section 4 (referred to as *basic Maude-NPA* to distinguish it from Maude-NPA with composition). In Section 5 we describe the syntax for sequential protocol composition and its operational semantics. In Section 6 we describe an implementation of composition via synchronization messages, and show that it is sound and complete with respect to the semantics given in Section 5. A presentation of the protocol transformation approach to protocol composition is described in Section 7, preparatory to an experimental evaluation and performance comparison between the two approaches given in Section 8. Finally, in Section 9 we conclude the paper and discuss related and future work, as well as lessons learned.

2. Motivating Examples

In this section we provide several motivating examples of sequential composition. These examples give a flavor for the variants of sequential composition that are used in constructing cryptographic protocols. A single parent protocol instance can be composed with either many instances of a child protocol, or with only one such child instance. Likewise, parent protocol roles can determine child protocol roles, or child protocol roles can be unconstrained. In Section 2.1 we provide an example of a one-parent, one-child protocol composition, which appeared in [24] and which is subject to a distance hijacking attack previously described in [18]; we also provide a corrected version that it is proved to be secure against distance hijacking. In Section 2.2 we provide an example of a one-parent, many-children protocol composition which is proved secure by our tool.

2.1. NSL Distance Bounding Protocol

In this example of a one-parent, one-child protocol composition, appeared in [24], the participants first use NSL to agree on a secret nonce. We reproduce the NSL protocol below using textbook Alice-and-Bob notation where $A \rightarrow B : m$ means participant with name A sending the message m to the participant with name B:

1. $A \rightarrow B : \{N_A, A\}_{pub(B)}$
2. $B \rightarrow A : \{N_A, N_B, B\}_{pub(A)}$
3. $A \rightarrow B : \{N_B\}_{pub(B)}$

where $\{M\}_{pub(A)}$ means message M encrypted using the public key of principal with name A, N_A and N_B are nonces generated by the respective principals, and we use the comma as message concatenation.

The agreed nonce N_A is then used in a distance bounding protocol described below. This is a type of protocol, originally proposed by [11] for smart cards, which has received new interest in recent years for its possible application in wireless environments [5]. The idea behind the protocol is that Bob uses the round trip time of a challenge-response protocol with Alice to compute an upper bound on her distance from him according to the following protocol:

4. $B \to A : N'_B$
   Bob records the time at which he sent N'_B
5. $A \to B : N_A \oplus N'_B$
   Bob records the time he receives the response and checks the equivalence $N_A = N_A \oplus N'_B \oplus N'_B$. If this holds, he uses the round-trip time of his challenge and response to estimate his distance from Alice

where $\oplus$ is the exclusive-or operator satisfying associativity (i.e., $X \oplus (Y \oplus Z) = (X \oplus Y) \oplus Z$) and commutativity (i.e., $X \oplus Y = Y \oplus X$) plus the self-cancellation property $X \oplus X = 0$ and the identity property $X \oplus 0 = X$. Note that Bob is the initiator and Alice is the responder of the distance bounding protocol, in contrast to the NSL protocol.

This protocol must satisfy two requirements. The first is that it must guarantee that $N_A \oplus N'_B$ was sent after N'_B was received, or Alice will be able to pretend that she is closer than she is. Note that if Alice and Bob do not agree on N_A beforehand, then Alice will be able to mount the following attack: $B \to A : N'_B$ and then $A \to B : N$. Of course, $N = N'_B \oplus X$ for some X. But Bob has no way of telling if Alice computed N using N'_B and X, or if she just sent a random N. Using NSL to agree on a $X = N_A$ in advance prevents this type of attack.

Bob also needs to know that the response comes from whom it is supposed to be from. In particular, an attacker should not be able to impersonate Alice. Using NSL to agree on N_A guarantees that only Alice and Bob can know N_A, so the attacker cannot impersonate Alice. However, it should also be the case that an attacker cannot pass off Alice's response as his own. This is not the case for the NSL distance bounding protocol, which is subject to a form of what has come to be known as the Distance Hijacking Attack [8] [1]. This attack was found by the authors of this paper by inspection and has been previously described in [18].

a) Intruder I runs an instance of NSL with Alice as the initiator and I as the responder, obtaining a nonce N_A.
b) I then runs an instance of NSL with Bob with I as the initiator and Bob as the responder, using N_A as the initiator nonce.
c) $B \to I : N'_B$ where I does not respond, but Alice, seeing this, thinks it is for her.
d) $A \to I : N'_B \oplus N_A$ where Bob, seeing this thinks this is I's response.

If Alice is closer to Bob than I is, then I can use this attack to appear closer to Bob than he is. This attack is a textbook example of a composition failure. NSL has all the properties of a good key distribution protocol, but fails to provide all the guarantees that are needed by the distance bounding protocol. However, in this case we can fix the problem, not by changing NSL, but by changing the distance bounding protocol so that it provides a stronger guarantee:

4. $B \to A : N'_B$

[1] This is not meant as a denigration of [24], whose main focus is on timing models in strand spaces, not the design of distance bounding protocols.

5. $A \rightarrow B : h(N_A, A) \oplus N'_B$ where h is a collision-resistant hash function.

As we show in our analysis in Section 8, this prevents the attack. I cannot pass off Alice's nonce as his own because it is now bound to her name.

The distance bounding example is a case of a one parent, one child protocol composition. Each instance of the parent NSL protocol can have only one child distance bounding protocol, since the distance bounding protocol depends upon the assumption that N_A is known only by A and B. But since the distance bounding protocol reveals N_A, it cannot be used with the same N_A more than once.

2.2. NSL Key Distribution Protocol

Our next example is a one parent, many children protocol composition, also using NSL. This type of composition arises, for example, in key distribution protocols in which the parent protocol is used to generate a master key, and the child protocol is used to generate a session key. In this case, one wants to be able to run an arbitrary number of instances of the child protocol with the same master key.

In the distance bounding example the initiator of the distance bounding protocol was always the child of the responder of the NSL protocol and vice versa. In the key distribution example, the initiator of the session key protocol can be the child of either the initiator or the responder of the NSL protocol. So, we have two possible child executions after NSL:

4. $A \rightarrow B : \{Sk_A\}_{h(N_A,N_B)}$
5. $B \rightarrow A : \{Sk_A; N'_B\}_{h(N_A,N_B)}$
6. $A \rightarrow B : \{N'_B\}_{h(N_A,N_B)}$

4. $B \rightarrow A : \{Sk_B\}_{h(N_A,N_B)}$
5. $A \rightarrow B : \{Sk_B; N'_A\}_{h(N_A,N_B)}$
6. $B \rightarrow A : \{N'_A\}_{h(N_A,N_B)}$

where Sk_A is the session key generated by principal A and h is again a collision-resistant hash function. This protocol is proved secure by our tool in Section 8.

3. Background on Term Rewriting

In this section we provide background on the concepts from term rewriting used in this paper. Due to space constraints, this section is rather terse and mainly intended to reference purposes. The reader should consult it as needed. Readers familiar with such terminology and notation can skip this section and proceed to the next section, where we provide examples of protocol specification.

We follow the classical notation and terminology from [36] for term rewriting and from [31,32] for rewriting logic and order-sorted notions.

We assume an *order-sorted signature* Σ with a finite poset of sorts $(\mathsf{S}, \leq)$ and a finite number of function symbols. We assume an S-sorted family $\mathcal{X} = \{\mathcal{X}_\mathsf{s}\}_{\mathsf{s}\in\mathsf{S}}$ of disjoint variable sets with each $\mathcal{X}_\mathsf{s}$ countably infinite. $\mathcal{T}_\Sigma(\mathcal{X})_\mathsf{s}$ denotes the set of terms of sort s, and $\mathcal{T}_{\Sigma,\mathsf{s}}$ the set of ground terms of sort s. We write $\mathcal{T}_\Sigma(\mathcal{X})$ and $\mathcal{T}_\Sigma$ for the corresponding term algebras. We write $\mathcal{V}ar(t)$ for the set of variables present in a term t. The set of positions of a term t is written $Pos(t)$, and the set of non-variable positions $Pos_\Sigma(t)$. The subterm of t at position p is $t|_p$, and $t[u]_p$ is the result of replacing $t|_p$ by u in t. In Maude-NPA, we use sorts to indicate such things as which terms are intended to be parts of messages, and which terms, such as strands, are part of the higher-level infrastructure. We also use sorts to provide restrictions on how messages may be constructed. For example, we can specify an encryption function as symbol e of arity two, where the first argument must be of sort key, while the second argument is of sort message, where $\mathsf{key} < \mathsf{message}$.

A *substitution* σ is a sort-preserving mapping from a finite subset of $\mathcal{X}$ to $\mathcal{T}_\Sigma(\mathcal{X})$. The set of variables assigned by σ is $Dom(\sigma)$ and the set of variables introduced by σ is $Ran(\sigma)$. The identity substitution is *id*. Substitutions are homomorphically extended to $\mathcal{T}_\Sigma(\mathcal{X})$. Application of substitution σ to term t is denoted by $t\sigma$. Thus $e(K,X)(\sigma = \{K \mapsto key(A,B), X \mapsto n(A,r)\} = e(key(A,B), n(A,r))$. The restriction of σ to a set of variables V is $\sigma|_V$. The composition of two substitutions is $x(\sigma\theta) = (x\sigma)\theta$ for $x \in \mathcal{X}$.

A Σ-*equation* is an unoriented pair $t = t'$, where $t \in \mathcal{T}_\Sigma(\mathcal{X})_\mathsf{s}$, $t' \in \mathcal{T}_\Sigma(\mathcal{X})_{\mathsf{s}'}$, and s and s' are sorts in the same connected component of the poset $(\mathsf{S}, \leq)$. Given a set E of Σ-equations, order-sorted equational logic induces a congruence relation $=_E$ on terms $t, t' \in \mathcal{T}_\Sigma(\mathcal{X})$; see [32]. Throughout this paper we assume that $\mathcal{T}_{\Sigma,\mathsf{s}} \neq \emptyset$ for every sort s. We denote the E-equivalence class of a term $t \in \mathcal{T}_\Sigma(\mathcal{X})$ as $[t]_E$ and the S-sorted families of sets of E-equivalence classes of all terms $\mathcal{T}_\Sigma(\mathcal{X})$ and $\mathcal{T}_\Sigma(\mathcal{X})_\mathsf{s}$ as $\mathcal{T}_{\Sigma/E}(\mathcal{X})$, and $\mathcal{T}_{\Sigma/E}(\mathcal{X})_\mathsf{s}$ for the quotient set of sort s, respectively. A substitution σ is *more general modulo* E than another substitution θ, written $\sigma \sqsupseteq_E \theta$, iff there is a substitution ρ such that $\sigma\rho =_E \theta$, i.e., such that $x\sigma\rho =_E x\theta$ for each $x \in \mathcal{X}$. In Maude-NPA we use equations to represent the properties of crypto systems. Thus, if we want to represent the fact that decryption with a key cancels out encryption with the same key, we can use the equation $d(K, e(K, X)) = X$.

For a set E of Σ-equations, an E-*unifier* for a Σ-equation $t = t'$ is a substitution σ s.t. $t\sigma =_E t'\sigma$. For $\mathcal{V}ar(t) \cup \mathcal{V}ar(t') \subseteq W$, a set of substitutions $\mathit{CSU}_E^W(t = t')$ is said to be a *complete* set of of E-unifiers of an equation $t = t'$ away from W iff: (i) each $\sigma \in \mathit{CSU}_E^W(t = t')$ is an E-unifier of $t = t'$; (ii) for any E-unifier ρ of $t = t'$ there is a $\sigma \in \mathit{CSU}_E^W(t = t')$ such that $\sigma|_W \sqsupseteq_E \rho|_W$; (iii) for all $\sigma \in \mathit{CSU}_E^W(t = t')$, $Dom(\sigma) \subseteq (\mathcal{V}ar(t) \cup \mathcal{V}ar(t'))$ and $Ran(\sigma) \cap W = \emptyset$. If the set of variables W is irrelevant or is understood from the context, we write $\mathit{CSU}_E(t = t')$ instead of $\mathit{CSU}_E^W(t = t')$. We say that E-unification is *finitary* if $\mathit{CSU}_E(t = t')$ contains a finite number of E-unifiers for any equation $t = t'$, and *unitary* if it contains most one. For example, E-unification when $E = \{d(K, e(K, X)) = X\}$ is finitary but not unitary. For example, the complete set of unifiers $\mathit{CSU}_E(d(K, X) = Y)$ contains two substitutions: $\sigma_1 = \{Y \mapsto d(K, X)\}$ and $\sigma_2 = \{X \mapsto e(K, Y)\}$.

A *rewrite rule* is an oriented pair $l \to r$, where $l \notin \mathcal{X}$ and $l, r \in \mathcal{T}_\Sigma(\mathcal{X})_\mathsf{s}$ for some sort $\mathsf{s} \in \mathsf{S}$. An *(unconditional) order-sorted rewrite theory* is a triple $\mathcal{R} = (\Sigma, E, R)$ with Σ an order-sorted signature, E a set of Σ-equations, and R a set of rewrite rules. A *topmost rewrite theory* (Σ, E, R) is a rewrite theory s.t. for each $l \to r \in R$, $l, r \in \mathcal{T}_\Sigma(\mathcal{X})_{\mathsf{State}}$ for a top sort State, and no operator in Σ has State as an argument sort. In Maude-NPA, topmost rewriting is used to describe which states can follow from other states. That is, these rewrite rules are topmost rules of form $S \to S'$, where S and S' are both of topmost terms sort State. The theory E used by Maude-NPA describes the equational properties of the cryptosystem.

The rewriting relation $\to_R$ on $\mathcal{T}_\Sigma(\mathcal{X})$ is $t \xrightarrow{p}_R t'$ (or $\to_R$) if $p \in Pos_\Sigma(t)$, $l \to r \in R$, $t|_p = l\sigma$, and $t' = t[r\sigma]_p$ for some σ. The relation $\to_{R/E}$ on $\mathcal{T}_\Sigma(\mathcal{X})$ is $=_E; \to_R; =_E$, i.e., $t \to_{R/E} s$ iff $\exists u_1, u_2 \in \mathcal{T}_\Sigma(\mathcal{X})$ s.t. $t =_E u_1 \to_R u_2 =_E s$. Note that $\to_{R/E}$ on $\mathcal{T}_\Sigma(\mathcal{X})$ induces a relation $\to_{R/E}$ on $\mathcal{T}_{\Sigma/E}(\mathcal{X})$ by $[t]_E \to_{R/E} [t']_E$ iff $t \to_{R/E} t'$. The relation $\to_{R/E}$ is undecidable in general, since E-congruence classes can be arbitrarily large, and the simpler relation $\to_{R,E}$ is used. The rewriting relation $\to_{R,E}$ on $\mathcal{T}_\Sigma(\mathcal{X})$ is performed by applying narrowing to representatives of $t \xrightarrow{p}_{R,E} t'$ (or $\to_{R,E}$) if $p \in Pos_\Sigma(t)$, $l \to r \in R$, $t|_p =_E l\sigma$, and $t' = t[r\sigma]_p$ for some σ. The narrowing relation $\rightsquigarrow_R$ on $\mathcal{T}_\Sigma(\mathcal{X})$ is $t \overset{p}{\rightsquigarrow}_{\sigma,R} t'$ (or $\rightsquigarrow_{\sigma,R}, \rightsquigarrow_R$) if $p \in Pos_\Sigma(t)$, $l \to r \in R$, $\sigma \in \mathit{CSU}_\emptyset(t|_p = l)$, and $t' = \sigma(t[r]_p)$. Assuming that E has a finitary and complete unification algorithm, the narrowing relation $\rightsquigarrow_{R,E}$ on $\mathcal{T}_\Sigma(\mathcal{X})$ is $t \overset{p}{\rightsquigarrow}_{\sigma,R,E} t'$ (or $\rightsquigarrow_{\sigma,R,E}, \rightsquigarrow_{R,E}$) if $p \in Pos_\Sigma(t)$, $l \to r \in R$, $\sigma \in \mathit{CSU}_E(t|_p = l)$, and $t' = (t[r]_p)\sigma$.

Maude-NPA uses narrowing modulo E to perform state space exploration. Its use of topmost rewrite theories provides several advantages; see [37]: (i) the relation $\rightarrow_{R,E}$ achieves the same effect as the relation $\rightarrow_{R/E}$, and (ii) we obtain a completeness result between narrowing ($\rightsquigarrow_{R,E}$) and rewriting ($\rightarrow_{R/E}$), in the sense that a reachability problem has a solution iff narrowing can find an instance of it.

For equational theories E describing the properties of cryptosystem, Maude-NPA uses $E = E' \uplus Ax$ such that the equations E' oriented as left-to-right rules are confluent, coherent, and terminating modulo axioms Ax such as commutativity (C), associativity-commutativity (AC), or associativity-commutativity plus identity (ACU) of some function symbols. We also require axioms Ax to be *regular*, i.e., for each equation $l = r \in Ax$, $\mathcal{V}ar(l) = \mathcal{V}ar(r)$.

Note that axioms such as commutativity (C), associativity-commutativity (AC), or associativity-commutativity plus identity (ACU) are regular. Maude-NPA has both dedicated and generic algorithms for solving unification problems in such theories $E' \uplus Ax$ under appropriate conditions [19].

4. Basic Maude-NPA's Execution Model and Protocol Analysis

In this section we describe the core syntax and semantics of Maude-NPA as described in [17]. We refer to it here as *basic Maude-NPA* to distinguish it from Maude-NPA with composition. When we describe features that will be modified once composition is added, we refer explicitly to "basic Maude-NPA". When a feature is the same for both versions we simply say "Maude-NPA."

In Maude-NPA the behaviors of protocols is modeled using rewrite theories. Briefly, a protocol $\mathcal{P}$ is a set of strands. Each strand is either a *protocol strand* that describes the actions of a role played by an honest principal, or an *intruder strand* describing the ways in which the intruder can derive new messages, e.g. by generating nonces or applying functions symbols to messages it already knows. Although the two are conceptually different, the are processed the same way in basic Maude-NPA. Thus, given a protocol $\mathcal{P}$, its behavior in basic Maude-NPA is modeled by the rewrite theory $(\Sigma_{\mathcal{P}}, E_{\mathcal{P}}, \mathcal{R}_{\mathcal{P}})$, where $\Sigma_{\mathcal{P}}$ is the signature defining the sorts and function symbols for the cryptographic functions and for all the state constructor symbols, $E_{\mathcal{P}}$ is a set of equations specifying the *algebraic properties* of the cryptographic functions and the state constructors, and $\mathcal{R}_{\mathcal{P}}$ is a set of rewrite rules representing the protocol's state changes. More specifically, given a protocol $\mathcal{P}$, a *state* in the protocol execution is an $E_{\mathcal{P}}$-equivalence class $[t]_{E_{\mathcal{P}}}$ with t a term of sort State, $[t]_{E_{\mathcal{P}}} \in T_{\Sigma_{\mathcal{P}}/E_{\mathcal{P}}}(\mathcal{X})_{\mathsf{State}}$. In basic Maude-NPA there are two types of algebraic properties: (i) *equational axioms*, such as commutativity, associativity-commutativity, or associativity-commutativity-identity, called *axioms*, and (ii) *equational rules*, called *equations*. Basic Maude-NPA includes two predefined sorts: (i) the sort Msg that allows the protocol specifier to describe other sorts as subsorts of the sort Msg, and (ii) the sort Fresh for representing fresh unguessable values, e.g., nonces.

Example 1 The specification of the NSL protocol in Maude-NPA is as follows. A nonce generated by principal A is denoted by $n(A, r)$, where r is a unique variable of sort Fresh and A denotes who generated the nonce. This representation makes it easier to specify and keep track of the origin of nonces. E.g., one can use the notation to specify a state in which a principal accepts a nonce as coming from A when it actually comes from some $B \neq A$. Concatenation of two messages, e.g., N_A and N_B, is denoted by the operator _;_, e.g., $n(A, r)\ ;\ n(B, r')$. Encryption of a message M with the public key of principal A is denoted by $pk(A, M)$, e.g., $\{N_B\}_{pub(B)}$ is denoted by $pk(B, n(B, r'))$. Encryption with the secret key of principal A is denoted by $sk(A, M)$. The signature Σ_{NSL} of the NSL protocol contains only terms such as $n(A, r)$, $M_1; M_2$, $pk(A, M)$, and $sk(A, M)$.

The equational theory of the NSL protocol contains no axioms and only the equations describing public/private encryption cancellation: $E_{NSL} = \{\, pk(A, sk(A, M)) = M,\ sk(A, pk(A, M)) = M\}$. ■

A protocol $\mathcal{P}$ is specified with a notation derived from strand spaces [20]. In a ***strand***, a local execution of a protocol by a principal is indicated by a sequence of messages $[msg_1^-,\ msg_2^+,\ msg_3^-,\ \ldots,\ msg_{k-1}^-,\ msg_k^+]$ where each msg_i is a term of sort Msg (i.e., $msg_i \in T_{\Sigma_\mathcal{P}}(\mathcal{X})_{\mathsf{Msg}}$). Strand items representing input messages are assigned a negative sign, and strand items representing output messages are assigned a positive sign. We write $m^\pm$ to denote m^+ or m^-, indistinctively. We often write $+(m)$ and $-(m)$ instead of m^+ and m^-, respectively. We make explicit the Fresh variables $r_1, \ldots, r_k (k \geq 0)$ generated by a strand by writing $:: r_1, \ldots, r_k :: [msg_1^\pm, \ldots, msg_n^\pm]$, where $r_1, \ldots, r_k$ appear somewhere in $msg_1^\pm, \ldots, msg_n^\pm$. Fresh variables generated by a strand are unique and this is enforced during execution. Furthermore, fresh variables are treated as constants that are never instantiated.

In Maude-NPA [16,17], strands evolve over time and thus we use the symbol | to divide past and future in a strand, i.e., $[nil, msg_1^\pm, \ldots, msg_{j-1}^\pm \mid msg_j^\pm, msg_{j+1}^\pm, \ldots, msg_k^\pm, nil]$, where $msg_1^\pm, \ldots, msg_{j-1}^\pm$ are the past messages, and $msg_j^\pm, msg_{j+1}^\pm, \ldots, msg_k^\pm$ are the future messages ($msg_j^\pm$ is the immediate future message). In this presentation we will often remove the nils to simplify the exposition, except when there is nothing else between the vertical bar and the beginning or end of a strand. If there is no risk of confusion, we may also remove the fresh variables appearing before the strand.

We write $\mathcal{P}$ for the set of strands in a protocol, including the strands that describe the intruder's behavior. When it is necessary to identify a strand $:: r_1, \ldots, r_k :: [msg_1^\pm, \ldots, msg_n^\pm]$ to distinguish it from other strands, we will do so via a ***role name*** in parentheses appearing before the strand, e.g. $(initiator) :: r_1, \ldots, r_k :: [msg_1^\pm, \ldots, msg_n^\pm]$.

Example 2 Let us continue Example 1. The two principal strands associated to the NSL protocol describing the three steps shown in Section 2.1 are as shown below.

$$:: r ::[nil \mid\ +(pk(B, n(A, r); A)), -(pk(A, n(A, r); N_B; B)), +(pk(B, N_B))]$$
$$:: r' ::[nil \mid\ -(pk(B, N_A; A)), +(pk(A, N_A; n(B, r'); B)), -(pk(B, n(B, r')))]$$

In the NSL protocol the intruder has the following capabilities: (i) it can perform encryption with any public key, (ii) it can only perform encryption with its own private key, (iii) it can concatenate two messages, and (iv) it can decompose a concatenation into each of its parts. For example, the intruder's ability to concatenate two messages M_1 and M_2 is described by the following strand:

$$:: nil ::[nil \mid\ -(M_1), -(M_2), +(M_1; M_2)]$$

■

A *state* in Maude-NPA is a pair consisting of a set of Maude-NPA strands and the *intruder knowledge* at that time. The set of Maude-NPA strands is unioned[2] together by an associative and commutativity union operator _&_ with identity operator $\emptyset$, along with an additional term describing the intruder knowledge

[2] In reality we consider a multiset of strands but duplicates are discarded as redundant, see [16,17].

at that point. The *intruder knowledge* is represented as a set of facts unioned[3] together with an associative and commutativity union operator $_,_$ with identity operator $\emptyset$. There are two kinds of intruder facts: *positive* knowledge facts (the intruder knows message m, i.e., $m{\in}\mathcal{I}$), and *negative* knowledge facts (the intruder does not yet know m but will know it in a future state, denoted by $m{\notin}\mathcal{I}$). We represent a state as a term

$$s_1 \& s_2 \& \cdots s_n \& (m_1{\in}\mathcal{I}, \ldots, m_k{\in}\mathcal{I}, m'_1{\notin}\mathcal{I}, \ldots, m'_j{\notin}\mathcal{I})$$

with $s_1 \& s_2 \& \cdots s_n$ the *set of strands* and $m_1{\in}\mathcal{I}, \ldots, m_k{\in}\mathcal{I}, m'_1{\notin}\mathcal{I}, \ldots, m'_j{\notin}\mathcal{I}$ the *intruder knowledge*, i.e., we consider the intruder knowledge as another state component, enclosed in parenthesis, to simplify the exposition.

We now describe the rewrite rules used in basic Maude-NPA to describe forward execution. When new strands are not added into the state, the rewrite rules $R_{\mathcal{P}}$ obtained from the protocol strands $\mathcal{P}$ are as follows, where L, L' are variables of the sort for lists of input and output messages ($+m$,$-m$), IK is a variable of the sort for sets of intruder facts ($m{\in}\mathcal{I}$,$m{\notin}\mathcal{I}$), SS is a variable of the sort for sets of strands, and M is a variable of sort Msg:

$$SS \,\&\, [L \mid M^-, L'] \,\&\, (M{\in}\mathcal{I}, IK) \rightarrow SS \,\&\, [L, M^- \mid L'] \,\&\, (M{\in}\mathcal{I}, IK) \tag{1}$$

$$SS \,\&\, [L \mid M^+, L'] \,\&\, IK \rightarrow SS \,\&\, [L, M^+ \mid L'] \,\&\, IK \tag{2}$$

$$SS \,\&\, [L \mid M^+, L'] \,\&\, (M{\notin}\mathcal{I}, IK) \rightarrow SS \,\&\, [L, M^+ \mid L'] \,\&\, (M{\in}\mathcal{I}, IK) \tag{3}$$

In a *forward execution* of the protocol strands, Rule (1) synchronizes an input message with a message already in the channel (i.e., learned by the intruder), Rule (2) accepts output messages but the intruder's knowledge is not increased, and Rule (3) accepts output messages and the intruder's knowledge is positively increased. Note that Rule (3) makes explicit *when* the intruder learned a message M, which is recorded in the previous state [4] by the negative fact $M{\notin}\mathcal{I}$.

New strands are added to the state by explicit introduction through dedicated rewrite rules (one for each honest or intruder strand). It is also the case that when we are performing a backwards search, only the strands that we are searching for are listed explicitly, and extra strands necessary to reach an initial state are dynamically added. Thus, when we want to introduce new strands into the explicit description of the state, we need to describe additional rules for doing that, as follows:

$$\text{For each } [\, l_1,\ u^+,\ l_2\,] \in \mathcal{P} : SS \ \&\, [\, l_1 \mid u^+, l_2\,] \,\&\, (u{\notin}\mathcal{I}, IK) \rightarrow SS \ \&\, (u{\in}\mathcal{I}, IK) \tag{4}$$

where u denotes a message, l_1, l_2 denote lists of input and output messages ($+m$,$-m$), IK denotes a variable of the sort for sets of intruder facts ($m{\in}\mathcal{I}$,$m{\notin}\mathcal{I}$), and SS denotes a variable of the sort for sets of strands.

Example 3 The rewrite rule introducing a new intruder strand during backwards execution associated to the concatenation of two learned messages is as follows:

[3] Again, in reality we consider a multiset of intruder facts but duplicates are discarded as redundant, see [16,17].

[4] Of course, in an actual forward execution of a protocol, the intruder knowledge only has positive facts. The usefulness of $m{\notin}\mathcal{I}$ becomes clear when we consider *backward* executions, so that at the beginning of a protocol execution all intruder knowledge will be "negative", i.e., to be learned in the future.

$$SS \,\&\, [M_1^-, M_2^- \mid (M_1; M_2)^+] \,\&\, ((M_1; M_2) \notin \mathcal{I}, IK) \rightarrow SS \,\&\, ((M_1; M_2) \in \mathcal{I}, IK)$$

■

In summary, for a protocol $\mathcal{P}$, the set of rewrite rules obtained from the protocol strands that are used for backwards narrowing reachability analysis *modulo* the equational properties $E_{\mathcal{P}}$ is $R_{\mathcal{P}} = \{(1), (2), (3)\} \cup \{(4)\}$.

An *initial state* is the final result of the backwards reachability process when an attack is found, and is described as follows:

1. in an initial state, all strands have the bar at the beginning, i.e., all strands are of the form $:: r_1, \ldots, r_j :: [\, nil \mid {m_1}^{\pm}, \ \ldots, \ {m_k}^{\pm} \,]$;
2. in an initial state, all the intruder knowledge is negative, i.e., all the items in the intruder knowledge are of the form $m \notin \mathcal{I}$ and therefore only to be known in the future.

From an initial state no further backwards reachability steps are possible.

Attack states describe not just single concrete attacks, but *attack patterns* (or if you prefer *attack scenarios*), which are specified symbolically as terms (with variables) whose instances are the final attack states we are looking for. Given an attack pattern, Maude-NPA tries to either find an instance of the attack or prove that no instance of such attack pattern is possible.

Example 4 In order to prove that the NSL protocol fixes the bug found in the Needham-Schroeder Public Key protocol (NSPK), i.e., the intruder cannot learn the nonce generated by Bob, we should specify the following attack state:

$$:: r :: [nil, -(pk(b, a; N_A)), +(pk(a, N_A; n(b, r); b)), -(pk(b, n(b, r))) \mid nil] \,\&\, (n(b, r)) \in \mathcal{I}$$

from which an initial state cannot be reached and has a finite search space, proving it secure. ■

5. Abstract Definition of Sequential Protocol Composition in Maude-NPA

Sequential composition of two protocols describes a situation in which one protocol (the *child*) can only execute after another protocol (the *parent*) has completed its execution, which allows the child protocol to use information generated during the execution of the parent protocol. The underlying idea of such a situation is that the end of the parent's protocol execution is *synchronized* with the beginning of the child's protocol execution. In this section we present a synchronization syntax and semantics which refines that of [18]. In Section 5.1 we first explain in detail the syntactic and semantics features necessary to express the synchronization among both protocols. Then, in Section 5.2 we provide an abstract definition of sequential composition of two or more protocols in Maude-NPA. Finally, in Section 5.4 we define a concrete execution model for the one-to-one and one-to-many protocol compositions by extending the basic Maude-NPA execution model. Throughout this paper, we will refer to the syntax and semantics explained in this section as *abstract composition syntax and semantics*.

5.1. Input/Output Parameters and Roles

In this section we describe in more detail the new features we need to make explicit in each protocol to later define abstract sequential protocol compositions. These features are identical to those defined in [18]. Each strand in a protocol specification in the Maude-NPA is now extended with *input* and *output parameters*. Input parameters are a sequence of variables of different sorts placed at the beginning of a strand. Output parameters are a sequence of terms placed at the end of a strand. The strand notation we will now use is $[\{\overrightarrow{I}\}, \overrightarrow{M}, \{\overrightarrow{O}\}]$ where $\overrightarrow{I}$ is a list of input parameter variables, $\overrightarrow{M}$ is a list of positive and negative terms in the strand notation of the Maude-NPA, and $\overrightarrow{O}$ is a list of output terms. Note that all the variables of $\overrightarrow{O}$ must appear in $\overrightarrow{M}$ or $\overrightarrow{I}$, i.e., no extra variables are allowed in sequential protocol composition outputs. The input and output parameters describe the exact assumptions about each principal. Note that we allow each honest or Dolev-Yao strand to be *labeled* (e.g. *NSL.init* or *NSL.resp*) to denote the "role" of that strand in the protocol, in contrast to the standard Maude-NPA syntax for strands. These strand labels play an important role in our protocol composition method.

Example 5 Following Examples 1 and 2, the protocol $\mathcal{P}$ with two strands associated to the three protocol steps shown in Section 2.1 is now described as follows:

$$
\begin{aligned}
(\textit{NSL.init}) :: r :: [\{A,B\}, &+(pk(B, n(A,r); A)),\\
&-(pk(A, n(A,r); N; B)),\\
&+(pk(B,N)),\\
&\{A, B, n(A,r), N\}].\\
(\textit{NSL.resp}) :: r :: [\{A,B\}, &-(pk(B, N; A)),\\
&+(pk(A, N; n(B,r); B)),\\
&-(pk(B, n(B,r))),\\
&\{A, B, N, n(B,r)\}].
\end{aligned}
$$

■

Example 6 Similarly to the NSL protocol, there are several technical details missing in the previous informal description of the Distance Bounding (DB) protocol. The exclusive-or operator is $\oplus$ and its equational properties are described using associativity and commutativity of $\oplus$ plus the equations[5] $X \oplus 0 = X$, $X \oplus X = 0$, and $X \oplus X \oplus Y = Y$. Since Maude-NPA does not yet include timestamps, we do not include all the actions relevant to calculating time intervals, sending timestamps, and checking them. The protocol $\mathcal{P}$ with two strands associated to the two protocol steps shown in Section 2.1 is described as follows:

[5]Note that the redundant equational property $X \oplus X \oplus Y = Y$ is necessary in Maude-NPA for coherence purposes; see [38,14].

$$
\begin{aligned}
(\textit{DB.init}) :: r :: [\{A, B, N_A\}, &+ (n(B, r)),\\
&- (n(B, r) \oplus N_A),\\
&\{A, B, N_A, n(B, r)\}].\\
(\textit{DB.resp}) :: nil :: [\{A, B, N_A\}, &- (N_B),\\
&+ (N_B \oplus N_A),\\
&\{A, B, N_A, N_B\}].
\end{aligned}
$$

This protocol specification makes clear that the nonce N_A used by the initiator is a parameter and is never generated by A during the run of DB. However, the initiator B does generate a new nonce. ■

Example 7 The previous informal description of the Key Distribution (KD) protocol also lacks several technical details, which we supply here. Encryption of a message M with key K is denoted by $e(K, M)$, e.g., $\{N'_B\}_{h(N_A,N_B)}$ is denoted by $e(h(n(A, r), n(B, r')), n(B, r''))$. Cancellation properties of encryption and decryption are described using the equations $e(X, d(X, Z)) = Z$ and $d(X, e(X, Z)) = Z$. Session keys are written $skey(A, r)$, where A is the principal's name and r is a Fresh variable. The protocol $\mathcal{P}$ with two strands associated to the KD protocol steps shown above is described as follows:

$$
\begin{aligned}
(\textit{KD.init}) :: r :: [\{A, B, K\}, &+ (e(K, skey(A, r)),\\
&- (e(K, skey(A, r); N)),\\
&+ (e(K, N)),\\
&\{A, B, K, skey(A, r), N\}].\\
(\textit{KD.resp}) :: r :: [\{A, B, K\}, &- (e(K, SK)),\\
&+ (e(K, SK; n(B, r))),\\
&- (e(K, n(B, r))),\\
&\{A, B, K, SK, n(B, r)\}].
\end{aligned}
$$

■

In the rest of this paper we remove irrelevant parameters (i.e. input parameters for strands with no parents, and output parameters for strands with no children) in order to simplify the exposition. Therefore, a strand is now a term of one of the following forms:

1. $[nil, \overrightarrow{M}, nil]$, i.e. a standard strand that cannot be connected to either a parent or a child strand,
2. $[\{\overrightarrow{I}\}, \overrightarrow{M}, nil]$, i.e. a *child strand* that can be connected to a parent strand,
3. $[nil, \overrightarrow{M}, \{\overrightarrow{O}\}]$, i.e. a *parent strand* that can be connected to a child strand,
4. $[\{\overrightarrow{I}\}, \overrightarrow{M}, \{\overrightarrow{O}\}]$, i.e. a strand that can be connected to both a parent and a child strand, or
5. $[\{\overrightarrow{I}\}, \{\overrightarrow{O}\}]$, i.e. a strand that can be connected to both a parent and a child strand, but without sending or receiving any message, called a *void strand*.

5.2. Strand and Protocol Composition

In this section we formally define sequential protocol composition in Maude-NPA. We first define the sequential composition of two strands, since this will help us to define sequential protocol composition in general. Intuitively, sequential composition of two strands describes a situation in which one strand (the *child*), can only execute after another strand (*the parent*) has completed its execution. Each composition of two strands is obtained by *matching the output parameters of the parent strand with the input parameters of the child strand* in a user-specified way. Note that it may be possible for a single parent strand to have more than one child strand.

Definition 1 (Sequential Strand Composition) *Given two strands* $(a) :: \overrightarrow{r_a} :: [\{\overrightarrow{I_a}\}, \overrightarrow{M_a}, \{\overrightarrow{O_a}\}]$ *and* $(b) :: \overrightarrow{r_b} :: [\{\overrightarrow{I_b}\}, \overrightarrow{M_b}, \{\overrightarrow{O_b}\}]$ *that are properly renamed to avoid variable sharing, a sequential strand composition is a triple of the form* $(a, b, MODE)$*, where* a *and* b *denote the parent and child roles, respectively, and MODE is either 1-1 or 1-*, indicating a one-to-one or one-to-many composition. This triple satisfies the following conditions for consistency:*

1. *both* $\overrightarrow{O_a}$ *and* $\overrightarrow{I_b}$ *have the same length, i.e.* $\overrightarrow{O_a} = m_1, \ldots, m_n$ *and* $\overrightarrow{I_b} = m'_1, \ldots, m'_n$*, and*
2. *there exists at least one substitution* σ *such that* $\overrightarrow{O_a} =_{E_\mathcal{P}} \overrightarrow{I_b}\sigma$*.*

We note that the definition of sequential strand composition given here differs from that given in [18] in that in Definition 1 each input parameter in a child strand is matched with the corresponding output parameters in the parent strand, while in [18] the user can choose which parameters are matched. This gives the user more flexibility, particularly in the case in which different children use different output parameters of the same parent. But it comes at the cost of being more complex to specify and implement. Moreover, the case of different children needing different output parameters can be taken care of by using "dummy" input parameters to match parental output parameters the child does not need, or more generally, by means of the protocol adapters described in Section 5.3.

Example 8 Let us consider again the NSL protocol of Example 5 and the DB protocol of Example 6. The composition of the NSL initiator strand and the DB responder strand is specified by the triple $(\textit{NSL.init}, DB.resp, \text{1-1})$. However, the NSL protocol had four output arguments while the DB protocol had three input arguments and we are required to adapt the syntax of the NSL protocol to have only the three arguments required by the DB protocol:

$$
\begin{aligned}
(\textit{NSL.init}) :: r :: [\{A, B\},\ & \\
& + (pk(B, n(A, r); A)), -(pk(A, n(A, r); N; B)), +(pk(B, N)), \\
& \{A, B, n(A, r)\}]. \\
(\textit{DB.resp}) :: nil :: [\{A, B, N_A\},\ & \\
& -(N_B), +(N_B \oplus N_A), \\
& \{A, B, N_A, N_B\}].
\end{aligned}
$$

■

Example 9 Let us consider again the NSL protocol of Example 5 and the KD protocol of Example 7. The composition of the NSL responder strand and the KD initiator strand is specified by the triple

(*NSL.resp*, *KD.init*,1-*). But again, the NSL protocol had different output arguments than the input arguments of the KD protocol and we are required to adapt the syntax of the NSL protocol as follows:

$$
\begin{aligned}
(\textit{NSL.resp}) :: r :: [\{A, B\},\ & \\
& -(pk(B, N; A)), +(pk(A, N; n(B, r); B)), -(pk(B, n(B, r))), \\
& \{B, A, h(N, n(B, r))\}].
\end{aligned}
$$

$$
\begin{aligned}
(\textit{KD.init}) :: r' :: [\{B, A, K\},\ & \\
& +(e(K, skey(B, r')), -(e(K, skey(B, r'); N')), +(e(K, N')), \\
& \{B, A, K, skey(B, r'), N'\}].
\end{aligned}
$$

such that the term $h(N, n(B, r))$ has the same sort as that of the input parameter K. ■

Intuitively, we can now define the sequential composition of two protocols as a set of sequential strand compositions.

Definition 2 (Sequential Composition of Two Protocols) *Given two protocols $\mathcal{P}_1$ and $\mathcal{P}_2$ that are properly renamed to avoid variable sharing, a sequential composition of both protocols, written $\mathcal{P}_1 ;_S \mathcal{P}_2$, is defined as a triple of the form $(\mathcal{P}_1, S, \mathcal{P}_2)$ where S denotes a set of strand compositions between a parent strand of $\mathcal{P}_1$ and a child strand of $\mathcal{P}_2$ of the form described in Definition 1. Note that the signature of such protocol composition is the union[6] of the signature of both protocols, i.e., $\Sigma_{\mathcal{P}_1;_S\mathcal{P}_2} = \Sigma_{\mathcal{P}_1} \cup \Sigma_{\mathcal{P}_2}$. Similarly, the set of equations specifying the algebraic properties of such protocol composition is the union[7] of the equations of both protocols, i.e., $E_{\mathcal{P}_1;_S\mathcal{P}_2} = E_{\mathcal{P}_1} \cup E_{\mathcal{P}_2}$.*

Example 10 Let us consider again both the NSL and DB protocols of Example 8 and their composition. The composition of both protocols, which is an example of a one-to-one composition, is specified as follows, indicating that the initiator of NSL can be composed with the responder of DB and the responder of NSL with the initiator of DB:

$$
\begin{aligned}
NSL ;_S\ DB = (NSL, \{&(NSL.init, DB.resp, 1{-}1), \\
&(NSL.resp, DB.init, 1{-}1)\}, DB)
\end{aligned}
$$

[6]Note that we allow shared items but require the user to solve any possible conflict. Operator and sort renaming is an option, as in the Maude module importation language, but we do not consider those details in this paper.

[7]We assume the combined equational theory satisfies all the requirements for having a finitary and complete unification algorithm.

The strands are left as follows, where we have removed irrelevant input and output parameters for clarity and simplicity:

$$\begin{aligned}
(\textit{NSL.init}) :: r :: [\,&+(pk(B, n(A,r); A)), -(pk(A, n(A,r); N; B)), +(pk(B,N)),\\
&\{A, B, n(A,r)\}]\\
(\textit{NSL.resp}) :: r :: [\,&-(pk(B, N; A)), +(pk(A, N; n(B,r); B)), -(pk(B, n(B,r))),\\
&\{A, B, N\}]\\
(\textit{DB.init}) :: r :: \ &[\{A, B, N_A\},\\
&+(n(B,r)), -(n(B,r) \oplus N_A)]\\
(\textit{DB.resp}) :: nil :: \ &[\{A, B, N_A\},\\
&-(N_B), +(N_B \oplus N_A)]
\end{aligned}$$

■

Example 11 Let us now consider the NSL and KD protocols of Example 9 and their composition. The composition of both protocols, which is an example of a one-to-many composition, is specified as follows, indicating that there are four possible compositions: the initiator of NSL composed with either the initiator or the responder of KD, and the responder of NSL composed with either the initiator or the responder of KD:

$$\begin{aligned}
NSL\ ;_S\ KD = (\ &\textit{NSL}, (\textit{NSL.init}, \textit{KD.init}, 1{-}*),\\
&(\textit{NSL.init}, \textit{KD.resp}, 1{-}*),\\
&(\textit{NSL.resp}, \textit{KD.init}, 1{-}*),\\
&(\textit{NSL.resp}, \textit{KD.resp}, 1{-}*)\}, \textit{KD})
\end{aligned}$$

The strands are as follows, where we have removed irrelevant input and output parameters for clarity and simplicity:

$$\begin{aligned}
(\textit{NSL.init}) :: r :: [\,&+(pk(B, n(A,r); A)), -(pk(A, n(A,r); N; B)), +(pk(B,N)),\\
&\{A, B, h(n(A,r), N)\}]\\
(\textit{NSL.resp}) :: r :: [\,&-(pk(B, N; A)), +(pk(A, N; n(B,r); B)), -(pk(B, n(B,r))),\\
&\{B, A, h(N, n(B,r))\}]\\
(\textit{KD.init}) :: r :: [\,&\{C, D, K\},\\
&+(e(K, skey(C,r)), -(e(K, skey(C,r); ND)), +(e(K, ND))]\\
(\textit{KD.resp}) :: r :: [\,&\{C, D, K\},\\
&-(e(K, SKD)), +(e(K, SKD; n(C,r)), -(e(K, n(C,r))]
\end{aligned}$$

Note that in the KD strands we use variables C and D to avoid confusion, since depending on how the NSL and KD protocols are composed, they will be instantiated as either the NSL initiator or the NSL responder name, represented by variables A and B, respectively. ■

In addition, we need to define the sequential composition of more than two protocols. Intuitively, the sequential composition of n protocols $\mathcal{P}_1, \ldots, \mathcal{P}_n$ is a sequence of two-protocol compositions, such that each protocol is composed with the previous protocol (except $\mathcal{P}_1$) and with the next protocol (except $\mathcal{P}_n$).

Definition 3 (Sequential Composition of n Protocols) *Given n protocols $\mathcal{P}_1, \ldots, \mathcal{P}_n$ that are properly renamed to avoid variable sharing, the sequential composition of all of them is denoted by:*

$$\mathcal{P}_1 ;_{S_1} \mathcal{P}_2 ;_{S_2} \mathcal{P}_3 ;_{S_3} \ldots ;_{S_{n-2}} \mathcal{P}_{n-1} ;_{S_{n-1}} \mathcal{P}_n$$

iff $\mathcal{P}_1 ;_{S_1} \mathcal{P}_2$, $\mathcal{P}_2 ;_{S_2} \mathcal{P}_3, \ldots, \mathcal{P}_{n-1} ;_{S_{n-1}} \mathcal{P}_n$ are sequential protocol compositions as explained in Definition 2.

5.3. Protocol Adapters

As we see from the examples in Section 5.2, putting the composition information inside the role specification itself instead of specifying them separately introduces a potential modularity issue if we want to reuse roles in different specifications, in that different compositions may require different information. For example, in one composition a child may require less information than a child in another composition with the same parent, as is the case in with NSL-DB versus NSL-KD. Or, it may be more convenient to present the information in different orders in either the parent or the child, as is the case for NSL-DB versus NSL-KD. Or, one child may need the result of applying a function to parent output, while the other may require the output without that function applied. Although some of these issues may be avoidable by careful planning, forcing the user to consider them in advance works against the sort of modularity we are trying to achieve.

As a solution to this problem we propose the use of *protocol adapters*, somewhat similar to the plug adapters one uses for overseas travel. A protocol adapter, applied to the output of a parent protocol, would perform the operations on it that would result in suitable input for the child protocol. Such operations would include, but would not necessarily be limited to:

1. restricting the output parameters to a subsequence used by a child;
2. permuting the output parameters in the order used by a child, and;
3. computing symbolic functions on the output.

In a similar way, the input parameters of a child protocol can be restricted to a subsequence or permuted to fit the output parameters of a parent. We note that it is currently possible to specify such role adapters directly from void strands, using a void strand that takes as its input the output parameters of the parent, and produces as its output the result of transforming these parameters into a format acceptable by the child. However, this is a suboptimal solution in that it introduces an extra narrowing step to address a purely syntactic issue. Thus, we are currently considering the best way of implementing protocol adapters on the syntactic level.

For each one-to-one strand composition $(a, b, 1{-}1)$ with
strand $(a)[\overrightarrow{M_a}, \{\overrightarrow{O_a}\}]$ for protocol $\mathcal{P}_1$, strand $(b)[\{\overrightarrow{I_b}\}, \overrightarrow{M_b}]$ for protocol $\mathcal{P}_2$,
and for each substitution σ s.t. $\overrightarrow{I_b}\sigma =_{E_\mathcal{P}} \overrightarrow{O_a}$, we add the following rules:

$$SS \,\&\, (a)\ [\overrightarrow{M_a} \mid \{\overrightarrow{O_a}\}] \,\&\, (b)\ [nil \mid \{\overrightarrow{I_b}\sigma\}, \overrightarrow{M_b}\sigma] \,\&\, IK$$
$$\rightarrow SS \,\&\, (a)\ [\overrightarrow{M_a}, \{\overrightarrow{O_a}\} \mid nil] \,\&\, (b)\ [\{\overrightarrow{I_b}\sigma\} \mid \overrightarrow{M_b}\sigma] \,\&\, IK \tag{5}$$

$$SS \,\&\, (a)\ [\overrightarrow{M_a} \mid \{\overrightarrow{O_a}\}] \,\&\, (b)\ [nil \mid \{\overrightarrow{I_b}\sigma\}, \overrightarrow{M_b}\sigma] \,\&\, IK$$
$$\rightarrow SS \,\&\, (b)\ [\{\overrightarrow{I_b}\sigma\} \mid \overrightarrow{M_b}\sigma] \,\&\, IK \tag{6}$$

Fig. 1. Forward semantics for one-to-one composition

5.4. Operational Semantics

As explained in Section 4, the operational semantics of protocol execution and analysis is based on rewrite rules denoting state transitions which are applied *modulo* the algebraic properties $E_\mathcal{P}$ of the given protocol $\mathcal{P}$. Therefore, in the one-to-one and one-to-many cases we must add new state transition rules in order to deal with protocol composition. Maude-NPA performs backwards search modulo $E_\mathcal{P}$ by reversing the transition rules expressed in a forward way; see Section 4.

In the one-to-one composition, we add the state transition rules of Figure 1 to the rewrite theory $(\Sigma_\mathcal{P}, E_\mathcal{P}, R_{B\mathcal{P}})$ of Section 4. Note that these transition rules are written in a forwards way but will be executed backwards, as the basic transition rules of Section 4. Rule 5 composes a parent and a child strand already present in the current state. Rule 6 is the same as Rule 5 but adds, in a backwards execution, a parent strand to the current state and composes it with an existing child strand. For example, given the composition of the NSL initiator's strand with the DB responder's strand $(\textit{NSL.init}, DB.resp, 1{-}1)$ where NSL.init and DB.resp were defined in Example 10, we add the following transition rule for Rule (5) where both the parent and the child strands are present and thus synchronized.

$$(\textit{NSL.init}) :: r ::$$
$$[\,+(pk(B, n(A,r); A)), -(pk(A, n(A,r); N; B)), +(pk(B, N)) \mid \{A, B, n(A,r)\}\,] \,\&$$
$$(\textit{DB.resp}) :: nil ::$$
$$[\,nil \mid \{A, B, n(A,r)\}, -(NB), +(NB * n(A,r))\,] \qquad \& \; SS \,\&\, IK$$
$$\longrightarrow$$
$$(\textit{NSL.init}) :: r ::$$
$$[\,+(pk(B, n(A,r); A)), -(pk(A, n(A,r); N; B)), +(pk(B, N)), \{A, B, n(A,r)\} \mid nil\,] \,\&$$
$$(\textit{DB.resp}) :: nil ::$$
$$[\,\{A, B, n(A,r)\} \mid -(NB), +(NB * n(A,r))\,] \qquad \& \; SS \,\&\, IK$$

For each one-to-many strand composition $(a, b, 1{-}*)$ with

strand $(a)[\overrightarrow{M_a}, \{\overrightarrow{O_a}\}]$ for protocol $\mathcal{P}_1$, strand $(b)[\{\overrightarrow{I_b}\}, \overrightarrow{M_b}]$ for protocol $\mathcal{P}_2$,

and for each substitution σ s.t. $\overrightarrow{I_b}\sigma =_{E_\mathcal{P}} \overrightarrow{O_a}$, we add one Rule 5, one Rule 6, and rule :

$$SS \,\&\, (a)\ [\overrightarrow{M_a} \mid \{\overrightarrow{O_a}\}] \,\&\, (b)\ [nil \mid \{\overrightarrow{I_b}\sigma\}, \overrightarrow{M_b}\sigma] \,\&\, IK$$
$$\rightarrow SS \,\&\, (a)\ [\overrightarrow{M_a} \mid \{\overrightarrow{O_a}\}] \,\&\, (b)\ [\{\overrightarrow{I_b}\sigma\} \mid \overrightarrow{M_b}\sigma] \,\&\, IK \tag{7}$$

Fig. 2. Forward semantics for one-to-many composition

One-to-many composition uses the rules in Figure 1 for the first child, plus an additional rule for subsequent children, described in Figure 2. Rule 7 composes a parent strand and a child strand but the bar in the parent strand is not moved, in order to allow further backwards child compositions. For example, given the composition of the NSL responder's strand with the KD initiator's strand $(\textit{NSL.resp}, KD.init, 1\text{-}*)$ where $NSL.resp$ and $KD.init$ are as defined in Example 11, we add the following transition rule for Rule (7):

$$\begin{aligned}
&(\textit{NSL.resp}) :: r ::\\
&[-(pk(B, NA; A)), +(pk(A, NA; n(B, r); B)), -(pk(B, n(B, r))) \mid \{B, A, h(NA, n(B, r))\}],\\
&(\textit{KD.init}) :: r' ::\\
&[nil \mid \{B, A, h(NA, n(B, r))\}, +(e(h(NA, n(B, r)), skey(B, r'))),\\
&-(e(h(NA, n(B, r)), skey(B, r'); N)), +(e(h(NA, n(B, r)), N))] \qquad \&SS \,\&\, IK\\
&\longrightarrow\\
&(\textit{NSL.resp}) :: r ::\\
&[-(pk(B, NA; A)), +(pk(A, NA; n(B, r); B)), -(pk(B, n(B, r))) \mid \{B, A, h(NA, n(B, r))\}],\\
&(\textit{KD.init}) :: r' ::\\
&[\{B, A, h(NA, n(B, r))\} \mid +(e(h(NA, n(B, r)), skey(B, r'))),\\
&-(e(h(NA, n(B, r)), skey(B, r'); N)), +(e(h(NA, n(B, r)), N))] \qquad \&SS \,\&\, IK
\end{aligned}$$

Thus, for a protocol composition $\mathcal{P}_1;_S\mathcal{P}_2$, the rewrite rules governing protocol execution are $R_{\mathcal{P}_1;_S\mathcal{P}_2} = \{(1), (2), (3)\} \cup (4) \cup (5) \cup (6) \cup (7)$. Note that the only generic rules are Rules (1),(2), (3) and all the other are obtained from the protocol specification, thus increasing the number of transition rules.

6. Composition via synchronization messages

In Section 5 we have provided an abstract syntax and a semantics for protocol composition, but this is not what has been implemented in the tool. There are two reasons for this, having to do with the fact that the rules in Figures 1 and 2 are parametrized by the strands in the two composed protocols. First of

all, this means that implementing the rules would require a significant modification of Maude-NPA to support the new composition data type. Secondly, the fact that each strand composition produces a new rule means that the number of rewrite rules is significantly increased. Increasing the number of rewrite rules can affect efficiency, since each rewrite rule must be tried at each narrowing step. Therefore, our approach has been to instead implement composition using communication between strands, which can be achieved using only slight modifications of constructs already present in Maude-NPA.

In [18] this communication was implemented via messages sent over the Dolev-Yao channel; this implementation, referred to as *synchronization by protocol transformation* has also been proved sound and complete in [35] with respect to the semantics given in Section 5. However, as we will show in Section 8 this had a serious impact on performance due to the interleaving of the additional Dolev-Yao messages, as well as making it more difficult to write specifications and attack states. Here, we present a modified version of Maude-NPA in which composition is achieved via *synchronization messages* that are passed *directly* between a parent and child strand without going through the Dolev-Yao channel. Although, as in the case of composition with respect to protocol transformation, it is necessary to add new rewrite rules, the rules are very similar to those of the basic Maude-NPA semantics, and require the addition of fewer parametrized rules than for protocol transformation. Composition of synchronization messages is still somewhat less expressive than the abstract semantics, in that the same role cannot engage in both one-to-one and one-to-many compositions. However, it can be proved sound and complete with respect to the abstract semantics with the same restrictions. We discuss how this apparent restriction can be mitigated in Section 6.1.

In Section 6.1 we introduce the notion of *synchronization* of protocol strands, a key idea underlying sequential protocol composition. In Section 6.2 we explain in detail the new Maude-NPA syntax for the specification of protocol composition via synchronization messages. Section 6.4 provides detailed information about the operational semantics of this direct implementation of protocol composition in Maude-NPA. Throughout this paper we will refer to these syntax and semantics as *synchronization via synchronization messages*. Finally, Section 6.5 proves the soundness and completeness of the semantics in Section 6.4 with respect to the abstract semantics in Section 5.4, thus proving that the semantics in Section 6.4 is a *correct* implementation of protocol composition in Maude-NPA. We use our two running examples (NSL-DB and NSL-KD) to illustrate our technique.

6.1. Synchronization Data Type Extension

As explained above, the underlying idea of a sequential protocol composition is that the end of the parent's protocol execution is *synchronized* with the beginning of the child's protocol execution. Since in Maude-NPA a protocol execution is denoted by a set of strands, we actually need to provide an infrastructure to express the notion of synchronization among strands, so that the strands of the parent protocol can in fact be "connected" with the strands of the child protocol.

Synchronization of strands can be achieved in Maude-NPA by extending its syntax to define a special type of message that we call *synchronization message*. The signature necessary to specify synchronization messages, written Σ_{Synch}, is as follows. Several sorts are added: Synch for the synchronization message, Role for user-definable constants denoting the roles in the protocol, RoleConnection for establishing which roles are the parent and which roles are the children, and Mode for choosing between one-to-one composition, denoted by constant 1-1, and one-to-many composition, denoted by 1-*. The synchronization messages are defined by patterns of the form:

$$\{a \rightarrow (b_1 b_2 \cdots b_j) \;;;\; \textit{Mode} \;;;\; \textit{Msg}\} \text{ and } \{(a_1 a_2 \cdots a_i) \rightarrow b \;;;\; \textit{Mode} \;;;\; \textit{Msg}\}.$$

The sort Role contains some constants defined by the user for role names $a_1 a_2 \cdots$, e.g. NSL.init or NSL.resp. The sort RoleConnection contains just one operator _→_, so that $a \rightarrow (b_1 b_2 \cdots b_j)$ specifies that a parent role a can have child roles b_1 through b_j, while $(a_1 a_2 \cdots a_i) \rightarrow b$ specifies the parent roles $a_1 a_2 \cdots a_i$ that a child b may have. Thus "NSL.init NSL.resp → KD.resp" indicates that either the initiator or the responder roles of the NSL protocol can be the parent of the responder role of the KD protocol. The information passed from parent to child is given in the third parameter, which is just a term of sort Msg, allowing the user to construct any message representing the information exchanged in the synchronization.

6.2. Syntax for Protocol Composition via synchronization messages

In this section we explain in detail how the Maude-NPA's syntax has been extended with synchronization messages (see Section 6.1) in order to support the input and output parameters of Section 5.1 and the abstract definition of protocol composition provided in Section 5.2. Synchronization messages are used to represent protocol compositions directly in the strand specification of the parent and child strands without any protocol transformation. A mapping from the notation for protocol composition of Section 5.2 into synchronization messages is described as follows.

Definition 4 (Parent Strand Synchronization) *Given two protocols $\mathcal{P}_1$ and $\mathcal{P}_2$, a set S of strand compositions, a role a of $\mathcal{P}_1$ of the form $(a)\ [\overrightarrow{M}, \{o_1, \ldots, o_n\}]$, and all the strand compositions for a in S, i.e., $(a, b_1, Mode), \ldots, (a, b_k, Mode)$, we define*

$$synch_S(a) = \left\{ (a)\ [\overrightarrow{M}, \{a \rightarrow b_1 \cdots b_k\ ;;\ \mathit{Mode}\ ;;\ (o_1; \cdots ; o_n)\}] \right\}$$

Definition 5 (Children Strand Synchronization) *Given two protocols $\mathcal{P}_1$ and $\mathcal{P}_2$, a set S of strand compositions, a role b of $\mathcal{P}_2$ of the form $(b)\ [\{i'_1, \ldots, i'_n\}, \overrightarrow{M'}]$, and all the strand compositions for b in S, i.e., $(a_1, b, Mode), \ldots, (a_k, b, Mode)$, we define*

$$synch_S(b) = \left\{ (b)\ [\{a_1 \cdots a_k \rightarrow b\ ;;\ \mathit{Mode}\ ;;\ (i'_1; \cdots ; i'_n)\}, \overrightarrow{M'}] \right\}$$

Definition 6 (Protocol Synchronization) *Given two protocols $\mathcal{P}_1$ and $\mathcal{P}_2$ that are properly renamed to avoid variable sharing, and a sequential protocol composition $\mathcal{P}_1 ;_S \mathcal{P}_2 = (\mathcal{P}_1, S, \mathcal{P}_2)$ where S denotes a set of strand compositions of the form $(a, b, MODE)$, the protocol synchronization, denoted $sync(\mathcal{P}_1 ;_S \mathcal{P}_2)$ is a single protocol which:*

1. *has signature $\Sigma_{\mathcal{P}_1} \cup \Sigma_{\mathcal{P}_2} \cup \Sigma_{Synch}$, where Σ_{Synch} is the new signature described in Section 6.1,*
2. *the equational theory is $E_{\mathcal{P}_1} \cup E_{\mathcal{P}_2}$*
3. *the set of strands is $synch(S)$, which is, by definition, the set of strands of the form $synch_S(r)$ for each role r in $\mathcal{P}_1$ and $\mathcal{P}_2$, and*
4. *all the protocol compositions of a role have the same mode (1-1 or 1-*), i.e., given a in $\mathcal{P}_1$ and $(a, b_1, MODE_1), \ldots, (a, b_k, MODE_k)$ in S, then $MODE_1 = \ldots = MODE_k$; similarly given b in $\mathcal{P}_2$ and $(a_1, b, MODE_1), \ldots, (a_k, b, MODE_k)$ in S, then $MODE_1 = \ldots = MODE_k$.*

As we shall see in Section 6.5, synchronization via synchronization messages implements our abstract composition semantics of Section 5.4, but in the next section we clarify the role connections of our framework.

6.3. Role Connections

As explained above, there are two types of synchronization messages

$$\{a \rightarrow (b_1 b_2 \cdots b_j)\ ;;\ \textit{Mode}\ ;;\ \textit{Msg}\} \text{ and } \{(a_1 a_2 \cdots a_i) \rightarrow b\ ;;\ \textit{Mode}\ ;;\ \textit{Msg}\}.$$

They correspond to two different parent or child situations associated to the abstract semantics that are now represented using synchronization messages.

First, in the abstract semantics there is nothing preventing a single instantiation of a parent role from having two or more children belonging to different roles, assuming both child roles are allowed by the specification. For example, in the NSL-KD composition we can have an instance of the NSL initiator strand being composed with both an instance of the KD initiator strand and an instance of the KD responder strand, since in Example 11, the NSL.init has output parameters $\{A, B, h(n(A, r), N)\}$ while both KD.init and KD.resp have input parameters $\{C, D, K\}$. Indeed, since we have one-to-many compositions for both KD.init and KD.resp, we could have an instance of NSL.init being composed with many different KD.init and many different KD.resp. In this case, we write "NSL.init $\rightarrow$ KD.init KD.resp" in the synchronization message of the parent strand NSL.init.

Second, in the abstract semantics we can have a child role that participates in multiple protocol compositions, though a single instantiation of a child role has only one parent. Again, in the NSL-KD composition we can have an instance of the KD initiator strand that can be synchronized with either an instance of the NSL initiator strand or the NSL responder strand, since in Example 11, the NSL initiator has output parameters $\{A, B, h(n(A, r), N)\}$, the NSL responder has output parameters $\{B, A, h(n(A, r), N)\}$, and the KD initiator has input parameters $\{C, D, K\}$. Note that, in contrast to one parent being composed with many child instances of the children roles, in this case an instance of a child role would be composed only with an instance of the parent role. In this case, we write "NSL.init NSL.resp $\rightarrow$ KD.init" in the synchronization message of the child strand KD.init.

We would like to stress that the restriction in Definition 6 about all the strand roles participating in composition using always the same mode has been the result of a conscious decision to trade off expressiveness against ease and readability of the specification. For example, we could have allowed roles to be used in one-to-one and one-to-many compositions by attaching modes to the names to each possible child role of a parent (and vice versa), e.g. "$a \rightarrow (b, 1-1),\ (c, 1-*)$ but decided that this complicates the specification too much. We note also that it is possible to simulate roles that compose with children (or parents) using different modes by using void strands. For example, instead of having a compose directly with b and c we could have a compose in $1-1$ mode with two void roles b_0 and c_0. Void role b_0 would then compose with child b in 1-1 mode and void role c_0 would compose with child c in $1-*$ model. The performance impact of the extra narrowing step introduced by the void role can be mitigated by the use of partial order reductions, as we do for other steps in which no messages are exchanged over the Dolev-Yao channel.

We note that it is not possible to simulate in the synchronization message syntax the case in which a single instantiation of a strand may have children in more than one mode, although this is possible in the abstract syntax. We believe that this is a reasonable price to pay. We may consider introducing this capability later, but if so it will be in a larger context in which we consider a much more expressive syntax and semantics that is given by the current abstract semantics. See Section 9.2 for a discussion.

In the following we provide the specification of our two examples of protocol composition, namely the NSL Distance Bounding protocol (NSL-DB) and the NSL Key Distribution protocol (NSL-KD),

presented in Sections 2.1 and 2.2, respectively, using the new synchronization message representation described above.

Example 12 We begin with our example of one-to-one protocol composition, i.e., the NSL-DB protocol. As explained in Section 2.1, the initiator of the DB protocol is always the child of the responder of the NSL protocol. The specification of the protocol strands using this syntax is as follows where the symbol $\oplus$ denotes the exclusive-or operator:

$$
\begin{aligned}
(\textit{NSL.init}) :: r :: [nil \mid &+(pk(B, n(A,r); A)), -(pk(A, n(A,r); NB; B)), +(pk(B, NB)),\\
&\{\textit{NSL.init} \rightarrow \textit{DB.resp} \,;;\, \text{1-1} \,;;\, (A; B; n(A,r))\}]\&\\
(\textit{NSL.resp}) :: r :: [nil \mid &-(pk(B, NA; A)), +(pk(A, NA; n(B,r); B)), -(pk(B, n(B,r))),\\
&\{\textit{NSL.resp} \rightarrow \textit{DB.init} \,;;\, \text{1-1} \,;;\, (A; B; NA)\}]\&\\
(\textit{DB.init}) :: r' :: [nil \mid &\{\textit{NSL.resp} \rightarrow \textit{DB.init} \,;;\, \text{1-1} \,;;\, (A; B; NA)\},\\
&+(n(B, r')), -(NA \oplus n(B, r'))]\&\\
(\textit{DB.resp}) :: nil :: [nil \mid &\{\textit{NSL.init} \rightarrow \textit{DB.resp} \,;;\, \text{1-1} \,;;\, (A; B; NA)\},\\
&-(N), +(NA \oplus N)]
\end{aligned}
$$

■

Example 13 Let us now continue with our example of a one-to-many protocol composition, i.e., the NSL-KD protocol. As explained in Section 2.2, the initiator of the session key protocol can be the child of either the initiator or responder of the NSL protocol. The specification of the strands of the NSL-KD protocol using the syntax for protocol composition via synchronization messages is as follows:

$$
\begin{aligned}
(\textit{NSL.init}) :: r :: [nil \mid &+(pk(B, n(A,r); A)), -(pk(A, n(A,r); NB; B)), +(pk(B, NB)),\\
&\{\textit{NSL.init} \rightarrow \textit{KD.init KD.resp} \,;;\, \text{1-*} \,;;\, (A; B; h(n(A,r), NB))\}\,]\&\\
(\textit{NSL.resp}) :: r :: [nil \mid &-(pk(B, NA; A)), +(pk(A, NA; n(B,r); B)), -(pk(B, n(B,r))),\\
&\{\textit{NSL.resp} \rightarrow \textit{KD.init KD.resp} \,;;\, \text{1-*} \,;;\, (B; A; h(NA, n(B,r)))\}\,]\&\\
(\textit{KD.init}) :: r' :: [nil \mid &\{\textit{NSL.init NSL.resp} \rightarrow \textit{KD.init} \,;;\, \text{1-*} \,;;\, (C; D; K)\},\\
&+(e(K, skey(C, r'))), -(e(K, skey(C, r'); N)), +(e(K, N))]\&\\
(\textit{KD.resp}) :: r' :: [nil \mid &\{\textit{NSL.init NSL.resp} \rightarrow \textit{KD.resp} \,;;\, \text{1-*} \,;;\, (C; D; K)\},\\
&-(e(K, SKD)), +(e(K, SKD; n(C, r'))), -(e(K, n(C, r')))]\&
\end{aligned}
$$

■

6.4. Operational Semantics of Composition via synchronization messages

In Section 5.4 we provided an operational semantics based on extra transition rules generated for each possible protocol composition and we differentiated between rules generated for one-to-one compositions and rules generated for one-to-many compositions. In this section we propose a simplified version of that operational semantics, which we call *composition via synchronization messages semantics*, that reduces the number of transition rules so that now we just have two *generic transition rules* and a set of *generated transition rules* for each strand in the same spirit of Rule (1) and Rules (4).

The two generic transition rules for protocol composition via synchronization messages are described in Figure 3. Note that these transition rules are written in a *forwards* way but will be executed backwards, as the basic transition rules of Section 4 and the abstract composition semantics of Section 5.4. The first generic transition Rule (8) is applicable to both one-to-one compositions and one-to-many compositions. This rule achieves the synchronization between both strands by means of the synchronization message. The second generic Rule (9) is applicable only to one-to-many compositions and represents the synchronization of a parent and a child without disabling the synchronization message of the parent.

These two generic rules synchronize an output parameter of an existing parent strand with an input message of an existing child strand. Both strands must be present in the state. The difference between a one-to-one and one-to-many composition is that the output parameter of the parent strand is kept in the same position of the parent strand for further synchronizations with other children strands.

As it happens in the basic Maude-NPA operational semantics of Section 4, we generate extra transitions rules from strands, in this case for protocol composition, as shown in Figure 4. Transition rules of the form (10), when executed backwards, allow adding to the state a new parent strand, whose output parameters will be synchronized with the input parameters of an already existing child strand. Note that the generated transitions rules (10) apply to *both* of the one-to-one or one-to-many composition cases. In each case, they describe a parent synchronizing with its first child.

For example, given the composition of the NSL initiator's strand and the DB responder's strand, where both strands were defined in Example 12, for Alice's strand

$$:: r :: [nil| + (pk(B, n(A, r); A)), -(pk(A, n(A, r); NB; B)), +(pk(B, NB)),$$
$$\{\textit{NSL.init} \rightarrow \textit{DB.resp} \; ; ; \; \text{1-1} \; ; ; \; (A; B; n(A, r))\}]$$

we add the following transition rule generated by Rule (10)

$$:: r :: [nil, +(pk(B, n(A, r); A)), -(pk(A, n(A, r); NB; B)), +(pk(B, NB)),$$
$$\mid \{\textit{NSL.init} \rightarrow \textit{DB.resp} \; ; ; \; \text{1-1} \; ; ; \; (A; B; n(A, r))\}] \;\&$$
$$:: RR :: [nil \mid \{\textit{NSL.init}\; R \rightarrow \textit{DB.resp} \; ; ; \text{1-1}; ; (A; B; n(A, r))\}, L] \qquad \& \; SS \; \& \; IK$$
$$\longrightarrow$$
$$:: RR :: [nil, \{\textit{NSL.init}\; R \rightarrow \textit{DB.resp} \; ; ; \; \text{1-1} \; ; ; \; (A; B; n(A, r))\} \mid L] \qquad \& \; SS \; \& \; IK$$

Thus, for a protocol composition $\mathcal{P}_1;_S \mathcal{P}_2$, the rewrite rules governing protocol execution in composition via synchronization messages are $R_{synch(\mathcal{P}_1;_S\mathcal{P}_2)} = \{(1), (2), (3)\} \cup (4) \cup \{(8), (9)\} \cup (10)$.

$$SS \,\&\, (\mathrm{a})[L \mid \{\mathrm{a} \to \mathrm{b\ R}\ ;;\ \mathrm{Mode}\ ;;\ \mathrm{M}\}] \,\&\, (\mathrm{b})[nil \mid \{\mathrm{a\ R'} \to \mathrm{b}\ ;;\ \mathrm{Mode}\ ;;\ \mathrm{M}\}, L'] \,\&\, IK$$
$$\to SS \,\&\, (\mathrm{a})[L, \{\mathrm{a} \to \mathrm{b\ R}\ ;;\ \mathrm{Mode}\ ;;\ \mathrm{M}\} \mid nil] \,\&\, (\mathrm{b})[\{\mathrm{a\ R'} \to \mathrm{b}\ ;;\ \mathrm{Mode}\ ;;\ \mathrm{M}\} \mid L'] \,\&\, IK \quad (8)$$

$$SS \,\&\, (\mathrm{a})[L \mid \{\mathrm{a} \to \mathrm{b\ R}\ ;;\ \texttt{1-*}\ ;;\ \mathrm{M}\}] \,\&\, (\mathrm{b})[nil \mid \{\mathrm{a\ R'} \to \mathrm{b}\ ;;\ \texttt{1-*}\ ;;\ \mathrm{M}\}, L'] \,\&\, IK$$
$$\to SS \,\&\, (\mathrm{a})[L \mid \{\mathrm{a} \to \mathrm{b\ R}\ ;;\ \texttt{1-*}\ ;;\ \mathrm{M}\}] \,\&\, (\mathrm{b})[\{\mathrm{a\ R'} \to \mathrm{b}\ ;;\ \texttt{1-*}\ ;;\ \mathrm{M}\} \mid L'] \,\&\, IK \quad (9)$$

where:

L, L' are variables of the sort for lists of input and output messages (+*m*,-*m*),

IK is a variable of the sort for sets of intruder facts $(m{\in}\mathcal{I}, m{\notin}\mathcal{I})$,

SS is a variable of the sort for sets of strands,

M is a variable of sort Msg,

a, b are variables of sort Role,

R, R' are variables denoting sets of roles, and

Mode is a variable of sort Mode

Fig. 3. Generic forward transition rules for composition via synchronization messages

For each strand definition $[\overrightarrow{M_a}, \{a \to b_1 \cdots b_k\ ;;\ \mathrm{mode}\ ;;\ \mathrm{msg}\}]$, and each $i \in \{1, \ldots, k\}$, we add a rule of the form:

$$SS \,\&\, (\mathrm{a})[\overrightarrow{M_a} \mid \{a \to b_1 \cdots b_k\ ;;\ \mathrm{mode}\ ;;\ \mathrm{msg}\}] \,\&$$
$$(b_i)[nil \mid \{a\ R \to b_i\ ;;\ \mathrm{mode}\ ;;\ \mathrm{msg}\}, L] \,\&\, IK$$
$$\to SS \,\&\, (b_i)[\{a\ R \to b_i\ ;;\ \mathrm{mode}\ ;;\ \mathrm{msg}\} \mid L] \,\&\, IK \quad (10)$$

where:

L is a variable of the sort for lists of input and output messages (+*m*,-*m*),

IK is a variable of the sort for sets of intruder facts $(m{\in}\mathcal{I}, m{\notin}\mathcal{I})$,

SS is a variable of the sort for sets of strands,

msg is a specific expression of sort Msg,

$a, b_1, \ldots, b_k$ are specific constants of sort Role,

R is a variable denoting sets of roles, and

mode is a specific constant of sort Mode

Fig. 4. Generated forward transition rules for composition via synchronization messages

$$trans_S(St) = \begin{cases} (b)[\{a_1 \cdots a_k \to b\,;;\, Mode\,;;\, I_b^\sharp\},\ \overrightarrow{b_1} \mid \overrightarrow{b_2}] \,\&\, St' & \text{if } (b)[\{\overrightarrow{I_b}\}, \overrightarrow{b_1} \mid \overrightarrow{b_2}] \in St, \\ & \quad (a_1, b, Mode), \ldots, (a_k, b, Mode) \in S, \\ & \quad trans_S(St - (b)) = St' \\ (b)[nil \mid \{a_1 \cdots a_k \to b\,;;\, Mode\,;;\, I_b^\sharp\},\ \overrightarrow{b_1}] \,\&\, St' & \text{if } (b)[nil \mid \{\overrightarrow{I_b}\}, \overrightarrow{b_1}] \in St, \\ & \quad (a_1, b, Mode), \ldots, (a_k, b, Mode) \in S, \\ & \quad trans_S(St - (b)) = St' \\ (a)[\overrightarrow{a_1} \mid \overrightarrow{a_2}, \{a \to b_1 \cdots b_k\,;;\, Mode\,;;\, O_a^\sharp\}] \,\&\, St' & \text{if } (a)[\overrightarrow{a_1} \mid \overrightarrow{a_2}, \{\overrightarrow{O_a}\}] \in St, \\ & \quad (a, b_1, Mode), \ldots, (a, b_k, Mode) \in S, \\ & \quad trans_S(St - (a)) = St' \\ (a)[\overrightarrow{a_1}, \{a \to b_1 \cdots b_k\,;;\, Mode\,;;\, O_a^\sharp\} \mid nil] \,\&\, St' & \text{if } (a)[\overrightarrow{a_1}, \{\overrightarrow{O_a}\} \mid nil] \in St, \\ & \quad (a, b_1, Mode), \ldots, (a, b_k, Mode) \in S, \\ & \quad trans_S(St - (a)) = St' \\ St & \text{otherwise} \end{cases}$$

where $I^\sharp$ (resp. $O^\sharp$) is equal to $\overrightarrow{I}$ (resp. $\overrightarrow{O}$) by replacing the comma "," by a semicolon ";" to denote concatenation of input and output parameters, e. g. input parameters $\overrightarrow{I} = \{A\,,\, B\,,\, NA\}$ is written as the sequence $I^\sharp = A\,;\, B\,;\, NA$.

Fig. 5. Function $trans$ between states valid according to the rewrite theory $\mathcal{R}_{\mathcal{P}_1\,;_S\mathcal{P}_2}$ and states valid according to the rewrite theory $\mathcal{R}_{synch(\mathcal{P}_1\,;_S\mathcal{P}_2)}$

Here, the reader can realize that this synchronization semantics for protocol composition contains two generic transition rules, Rules (8) and (9), and one transition rule for each protocol composition from Rule (10), whereas the protocol composition presented in Section 5 produces several transition rules for each protocol composition. Indeed, this simpler semantics for protocol composition requires fewer rules distinguishing one-to-one and one-to-many compositions than the abstract semantics.

6.5. Soundness and Completeness

In this section we prove soundness and completeness of the operational semantics composition via synchronization messages presented in Section 6.4 with respect to the abstract compositional operational semantics of Section 5.4 under the restriction that a each child role (respectively parent role) and compose with parent (respectively child roles) in at most one mode.

First, we must relate protocol states using the protocol composition rewrite rules of Section 5.4 and protocol states in the composition via synchronization messages. Throughout this section, when we can avoid confusion, a state St is called *valid according to a rewrite theory* $\mathcal{R}$ if it is a valid term of sort State with respect to the order-sorted signature of $\mathcal{R}$.

Definition 7 (Bijective function *trans*) *Let $\mathcal{P}_1$ and $\mathcal{P}_2$ be two protocols and $\mathcal{P}_1\,;_S\,\mathcal{P}_2$ their composition. Let $\mathcal{R}_{\mathcal{P}_1\,;_S\mathcal{P}_2}$ be the rewrite theory associated in Section 5.4 to the abstract protocol composition $\mathcal{P}_1\,;_S\,\mathcal{P}_2$ and $\mathcal{R}_{synch(\mathcal{P}_1\,;_S\mathcal{P}_2)}$ be the rewrite theory associated in Section 6.4 to composition via synchronization messages. We define the function $trans_S$ mapping states valid according to the rewrite theory $\mathcal{R}_{\mathcal{P}_1\,;_S\mathcal{P}_2}$ to states valid according to the rewrite theory $\mathcal{R}_{synch(\mathcal{P}_1\,;_S\mathcal{P}_2)}$ as specified in Figure 5, and its inverse function $trans_S^{-1}$ as specified in Figure 6.*

The following auxiliary results ensure that there is an appropriate connection between states of both rewrite theories.

$$trans_S^{-1}(St) = \begin{cases} (b)[\{\overrightarrow{I_b}\}, \overrightarrow{b_1} \mid \overrightarrow{b_2}] \,\&\, St' & \text{if } (b)[\{a_1 \cdots a_k \to b\,;;\, Mode\,;;\, I_b^\sharp\},\ \overrightarrow{b_1} \mid \overrightarrow{b_2}] \in St, \\ & \quad trans_S^{-1}(St - (b)) = St' \\ (b)[nil \mid \{\overrightarrow{I_b}\}, \overrightarrow{b_1}] \,\&\, St' & \text{if } (b)[nil \mid \{a_1 \cdots a_k \to b\,;;\, Mode\,;;\, I_b^\sharp\},\ \overrightarrow{b_1}] \in St, \\ & \quad trans_S^{-1}(St - (b)) = St' \\ (a)[\overrightarrow{a_1} \mid \overrightarrow{a_2}, \{\overrightarrow{O_a}\}] \,\&\, St' & \text{if } (a)[\overrightarrow{a_1} \mid \overrightarrow{a_2}, \{a \to b_1 \cdots b_k\,;;\, Mode\,;;\, O_a^\sharp\}] \in St, \\ & \quad trans_S^{-1}(St - (a)) = St' \\ (a)[\overrightarrow{a_1}, \{\overrightarrow{O_a}\} \mid nil] \,\&\, St' & \text{if } (a)[\overrightarrow{a_1}, \{a \to b_1 \cdots b_k\,;;\, Mode\,;;\, O_a^\sharp\} \mid nil] \in St, \\ & \quad trans_S^{-1}(St - (a)) = St' \\ St & \text{otherwise} \end{cases}$$

where $I^\sharp$ (resp. $O^\sharp$) is equal to $\overrightarrow{I}$ (resp. $\overrightarrow{O}$) by replacing the comma "," by a semicolon ";" to denote concatenation of input and output parameters, e. g. input parameters $\overrightarrow{I} = \{A\,,\ B\,,\ NA\}$ is written as the sequence $I^\sharp = A\,;\ B\,;\ NA$.

Fig. 6. Function $trans^{-1}$ between states valid according to the rewrite theory $\mathcal{R}_{synch(\mathcal{P}_1\,;_S\mathcal{P}_2)}$ and states valid according to the rewrite theory $\mathcal{R}_{\mathcal{P}_1\,;_S\mathcal{P}_2}$

Lemma 1 *Let $\mathcal{P}_1$ and $\mathcal{P}_2$ be two protocols and $\mathcal{P}_1\,;_S\,\mathcal{P}_2$ their composition. Let $\mathcal{R}_{\mathcal{P}_1\,;_S\mathcal{P}_2}$ be the rewrite theory associated in Section 5.4 to the protocol composition $\mathcal{P}_1\,;_S\,\mathcal{P}_2$ and $\mathcal{R}_{synch(\mathcal{P}_1\,;_S\mathcal{P}_2)}$ be the rewrite theory associated in Section 6.4 to the composition via synchronization messages.*

Then $trans_S$ defined in Definition 7 is a bijective function from terms of sort State *in $\mathcal{R}_{synch(\mathcal{P}_1\,;_S\mathcal{P}_2)}$ to terms of sort* State *in $\mathcal{R}_{\mathcal{P}_1\,;_S\mathcal{P}_2}$, and has $trans_S^{-1}$ as its inverse function.*

Proof. By structural induction on the functions $trans_S$ and $trans_S^{-1}$ given in Figures 5 and 6. The base case is a state St that has no strand with input or output parameters, since $trans_S(St) = St$. For the inductive case we consider only the case when St contains a strand of the form $(b)[\{\overrightarrow{I_b}\}, \overrightarrow{b_1} \mid \overrightarrow{b_2}]$ and all the other cases are similar. Let $St = (b)[\{i_1, \ldots, i_n\}, m_1^\pm, \ldots, m_i^\pm \mid m_{i+1}^\pm, \ldots, m_k^\pm] \& ss \& ik$ where ss denotes a set of strand instances and ik the intruder knowledge of the state. Let $(a_1, b, mode), \ldots, (a_k, b, mode)$ be all the composition triples in S involving role b. By induction hypothesis we have that $trans_S^{-1}(trans_S(ss \& ik)) = ss \& ik$. Then, by applying function $trans$ to St we have that the strand instance b is transformed into $(b)[\{a_1 \cdots a_k \to b\,;;\, mode\,;;\, (i_1\,; \cdots\,; i_n)\},\ m_1^\pm, \ldots, m_i^\pm \mid m_{i+1}^\pm, \ldots, m_k^\pm]$. But then it is easy to see that when we apply $trans^{-1}$ to this transformed strand, we simply remove the synchronization message and get the same strand instance b. Therefore, $trans_S^{-1}(trans_S(St)) = St$. □

Let us now relate backwards narrowing steps using the rewrite theory associated to the composition via synchronization messages of Section 6.4 (i.e., $\mathcal{R}_{synch(\mathcal{P}_1\,;_S\mathcal{P}_2)}$) w.r.t. backwards narrowing using the rewrite theory associated to the abstract protocol composition of Section 5.4 (i.e., $\mathcal{R}_{\mathcal{P}_1\,;_S\mathcal{P}_2}$). Note that in this case a backwards narrowing step performed with a rule of $\mathcal{R}_{\mathcal{P}_1\,;_S\mathcal{P}_2}$ always corresponds to one backwards narrowing step with a rule of $\mathcal{R}_{\mathcal{P}_1\,;_S\mathcal{P}_2}$, since no extra messages are introduced to synchronize parent and child strands.

Lemma 2 (Bisimulation) *Let $\mathcal{P}_1$ and $\mathcal{P}_2$ be two protocols and $\mathcal{P}_1\,;_S\,\mathcal{P}_2$ their composition. Let $\mathcal{R}_{\mathcal{P}_1\,;_S\mathcal{P}_2}$ be the rewrite theory associated in Section 5.4 to the abstract protocol composition $\mathcal{P}_1\,;_S\,\mathcal{P}_2$, and $\mathcal{R}_{synch(\mathcal{P}_1\,;_S\mathcal{P}_2)}$ be the rewrite theory associated in Section 6.4 to composition via synchronization messages.*

Given two states St_1 and St_2 valid according to the rewrite theory $\mathcal{R}_{\mathcal{P}_1 ;_S \mathcal{P}_2}$ such that trans$_S(St_1) = St'_1$, trans$_S(St_2) = St'_2$, $St_1 \overset{1}{\rightsquigarrow}_{\rho, \mathcal{R}^{-1}_{\mathcal{P}_1 ;_S \mathcal{P}_2}, E_{\mathcal{P}_1 ;_S \mathcal{P}_2}} St_2$ iff $St'_1 \overset{1}{\rightsquigarrow}_{\sigma, \mathcal{R}^{-1}_{synch(\mathcal{P}_1 ;_S \mathcal{P}_2)}, E_{\mathcal{P}_1 ;_S \mathcal{P}_2}} St'_2$.

Proof. We prove the result by case analysis on the applicable rewrite rules. First, let us recall the different rules that are applicable: for a term St_1 valid according to the rewrite theory $\mathcal{R}_{\mathcal{P}_1 ;_S \mathcal{P}_2}$ we can apply the reversed version of Rules (1), (2), and (3) plus the reversed version of rules in any of the sets (4), (5), (6), and (7), whereas for the term $St'_1 = trans_S(St_1)$ valid according to the rewrite theory $\mathcal{R}_{synch(\mathcal{P}_1 ;_S \mathcal{P}_2)}$ we can apply the reversed version of Rules (1), (2), (3), (8), and (9) plus the reversed version of rules in any of the sets (4) and (10). Second, we consider four possibilities below but only show in detail cases (a) and (b), since cases (c) and (d) are similar to case (b).

(a) When the Rules (1),(2), and (3), as well as any rule in the set (4), are applied, they do not involve any composition and, since the function *trans*$_S$ is a bijection, the same type of rule would be applicable to St'_1.

(b) A rule in the set (5) corresponds to an application of Rule (8) (synchronizing the input parameters of the child strand with the output parameters of the parent strand). In this case, the reversed version of a rule of the following form in set (5) has been applied to state St_1

$$\begin{aligned} & SS \,\&\, (a)\ [\overrightarrow{M_a} \mid \{\overrightarrow{O_a}\}] \,\&\, (b)\ [nil \mid \{\overrightarrow{I_b}\sigma\}, \overrightarrow{M_b}\sigma] \,\&\, IK \\ \rightarrow & SS \,\&\, (a)\ [\overrightarrow{M_a}, \{\overrightarrow{O_a}\} \mid nil] \,\&\, (b)\ [\{\overrightarrow{I_b}\sigma\} \mid \overrightarrow{M_b}\sigma] \,\&\, IK \end{aligned}$$

where $(a, b, 1{-}1) \in S$, $(a)[\overrightarrow{M_a}, \{\overrightarrow{O_a}\}]$ is a role in $\mathcal{P}_1$, $(b)[\{\overrightarrow{I_b}\}, \overrightarrow{M_b}]$ is a role in $\mathcal{P}_2$, $\overrightarrow{O_a}, \overrightarrow{I_b}$ are two sequences of terms with variables, $\overrightarrow{M_a}, \overrightarrow{M_b}$ are two sequences of input and output messages, $\overrightarrow{O_a} =_{E_\mathcal{P}} \overrightarrow{I_b}\sigma$, and only SS and IK are variables. Since this rule was applied, there is a substitution ρ such that $(a)\ [\overrightarrow{M_a}\rho, \{\overrightarrow{O_a\rho}\} \mid nil]$ and $(b)\ [\{\overrightarrow{I_b}\sigma\rho\} \mid \overrightarrow{M_b}\sigma\rho]$ are strand instances in St_1. But, by application of the *trans* function, there are strands $(a)\ [\overrightarrow{M_a}\rho, \{a \rightarrow b_1 \cdots b_{i-1}\ b\ b_i \cdots b_k; ; \text{1-1}; ; O^\sharp_a\rho\} \mid nil]$ and $(b)\ [\{a_1 \cdots a_{j-1}\ a\ a_j \cdots a_n \rightarrow b; ; \text{1-1}; ; I^\sharp_b\sigma\rho\} \mid \overrightarrow{M_b}\sigma\rho]$ in St'_1. Now, since $\overrightarrow{O_a}\rho =_{E_\mathcal{P}} \overrightarrow{I_b}\sigma\rho$, the reversed version of Rule (8) is applicable

$$\begin{aligned} & SS \,\&\, [L \mid \{\text{a} \rightarrow \text{b R} ;; \text{ Mode} ; ; \text{M}\}] \,\&\, [nil \mid \{\text{a R'} \rightarrow \text{b} ;; \text{ Mode} ; ; \text{M}\}, L'] \,\&\, IK \\ \rightarrow & SS \,\&\, [L, \{\text{a} \rightarrow \text{b R} ;; \text{ Mode} ; ; \text{M}\} \mid nil] \,\&\, [\{\text{a R'} \rightarrow \text{b} ;; \text{ Mode} ; ; \text{M}\} \mid L'] \,\&\, IK \end{aligned}$$

where SS, L, a, b, R, Mode, M, R', L', IK are variables.

(c) A rule in the set (6) corresponds to an application of a rule in the set (10) (introducing a new parent strand and composing it with an existing child strand).

(d) A rule in the set (7) corresponds to an application of Rule (9) (synchronizing the output parameters of the parent strand with the already accepted input parameters of the child strand, but without moving the bar in the parent strand). □

Finally, we can put everything together into the following result.

Theorem 1 (Soundness and Completeness) *Let $\mathcal{P}_1$ and $\mathcal{P}_2$ be two protocols and $\mathcal{P}_1 ;_S \mathcal{P}_2$ their composition, as defined in Section 5.2. Let $\mathcal{R}_{synch(\mathcal{P}_1 ;_S \mathcal{P}_2)}$ be the rewrite theory associated to composition via synchronization messages defined above in Section 6.4, and let $\mathcal{R}_{\mathcal{P}_1 ;_S \mathcal{P}_2}$ be the rewrite theory associated to the abstract protocol composition, as described in Section 5.4.*

Given a state St valid according to $\mathcal{R}_{\mathcal{P}_1 ;_S \mathcal{P}_2}$ and an initial state St_{ini} such that trans$(St) = St'$ and trans$(St_{ini}) = St'_{ini}$, then St_{ini} is reachable from St by backwards narrowing in $\mathcal{R}_{\mathcal{P}_1 ;_S \mathcal{P}_2}$ iff St'_{ini} is reachable from St' by backwards narrowing in $\mathcal{R}_{synch(\mathcal{P}_1 ;_S \mathcal{P}_2)}$.

Proof. By successive applications of Lemma 2. □

In the case of sequential composition of n protocols $\mathcal{P}_1 ;_{S_1} \mathcal{P}_2 ;_{S_2} \mathcal{P}_3 ;_{S_3} \ldots ;_{S_{n-2}} \mathcal{P}_{n-1} ;_{S_{n-1}} \mathcal{P}_n$ as described in Definition 3, we can define a function *trans*$_{S_1,S_2,S_3,\ldots,S_{n-1}}$ between states valid according to the rewrite theory $\mathcal{R}_{\mathcal{P}_1 ;_{S_1} \mathcal{P}_2 ;_{S_2} \mathcal{P}_3 ;_{S_3} \cdots ;_{S_{n-2}} \mathcal{P}_{n-1} ;_{S_{n-1}} \mathcal{P}_n}$ and states valid according to the rewrite theory $\mathcal{R}_{synch(\mathcal{P}_1 ;_{S_1} \mathcal{P}_2 ;_{S_2} \mathcal{P}_3 ;_{S_3} \cdots ;_{S_{n-2}} \mathcal{P}_{n-1} ;_{S_{n-1}} \mathcal{P}_n)}$ with the only requirement that the role names of a protocol $\mathcal{P}_i$ have to be different from the role names of all other protocols $\mathcal{P}_j$, $j \neq i$. This requirement ensures that each strand instance can be easily associated to one of the protocols; otherwise we may have a strand instance being associated to several protocol states. We are working on relaxing this condition, perhaps via use of role adapters (Section 5.3).

7. Composition via Protocol Transformation

In this section we describe our previous approach to composition using protocol transformation, presented in [18,34]. This section provides background for Section 8, in which the performance of composition via synchronization messages is compared with its predecessor.

In [18], we presented an approach for protocol composition where we defined a notion of sequential protocol composition slightly different from the one presented in Section 5 and the transition rules associated to such a composition. We did not implement those transition rules in the Maude-NPA. Instead, we defined a protocol transformation that achieved the same effect using the existing Maude-NPA tool. Proofs of soundness and completeness of the protocol transformation for the transition rules of [18] were provided in [34].

However, when experimenting with actual protocol composition examples, we realized that such a protocol composition and its semantics were quite complex and produced too many transition rules for a concrete protocol composition. This led us to refine such protocol composition and its transition rules in the considerably simpler form now presented in Section 5. Besides being simpler, it has also a more effective protocol composition semantics, more suitable for implementation. We then investigated two routes to obtaining a Maude-NPA implementation of the simpler composition notion and it semantics of Section 5:

1. the more direct route based on synchronization messages presented in Section 6.1; and
2. the older route from [18] based on *protocol transformation*, but now according to the new composition notion and associated semantics of Section 5.

This was then used as a basis to compare more carefully which of these two possible implementation routes would be the best. To begin with, we wanted to prove that both (1) and (2) above provided *correct* implementations. The correctness of the synchronization-based route of (1) has been proved in Section 6.5. Similarly, in analogy with [18,34], the redefined and adapted notion of protocol transforma-

tion in (2) has been proved *correct* in [35] with respect to the new protocol composition semantics of Section 5. Once we were sure that both implementation routes were correct, we proceeded to compare their ease of use, simplicity, and performance through concrete case studies.

The rest of this section briefly describes route (2), based on the protocol transformation. A more detailed comparison of ease of use, simplicity, and performance between (1) and (2) is postponed until Section 8.

Given two protocols $\mathcal{P}_1$ and $\mathcal{P}_2$, its sequential composition implemented via the redefined and adapted protocol transformation in (2), written $\Phi(\mathcal{P}_1 ;_S \mathcal{P}_2)$, is a single, composed protocol specification where:

1. Sorts, symbols, and equational properties of both protocols are put together into a single specification. As explained in Footnote 2 in Section 4, we allow shared items but require the user to solve any possible conflict.
2. A new sort Param is defined to denote input and output parameters. The sort Param is disjoint from the sort Msg used by the protocol in the honest and intruder strands to ensure that an intruder cannot fake a composition.
3. For each composition (a, b, MODE) with underlying substitution σ such that $\overrightarrow{O_a} =_{E_\mathcal{P}} \overrightarrow{I_b}\sigma$, we transform the input parameters $\{\overrightarrow{I_b}\}$ into an input message exchange of the form $-(\overrightarrow{I_b})$, and the output parameters $\{\overrightarrow{O_a}\}$ into an output message exchange of the form $+(\overrightarrow{I_b}\sigma)$. In order to avoid type conflicts, we use a *dot* for concatenation within protocol composition exchange messages, e.g. input parameters $\overrightarrow{I} = \{A, B, NA\}$ are transformed into the sequence $\dot{I} = A \,.\, B \,.\, NA$.
4. Each composition is uniquely identified by using a composition identifier (a variable of sort Fresh). Strands exchange such composition identifier by using input/output messages of the form $role_j(r)$, which make the role explicit. The sort Role of these messages is disjoint from the sorts Param and Msg.

   (a) In a one-to-one protocol composition, the child strand uniquely generates a fresh variable that is added to the area of fresh identifiers at the beginning of its strand specification. This fresh variable must be passed from the child to the parent before the parent generates its output parameters and sends them back again to the child. What this simulates in practice is the uniqueness of the one-to-one composition, since the parent can generate a single such message.

   (b) In a one-to-many protocol composition, the parent strand uniquely generates a fresh variable that is passed to each child. Since an (a priori) unbounded number of children will be composed with it, no reply to the fresh variable is expected by the parent from the children. Note that all the children strands receive the same fresh variables from the parent.

Let us illustrate this protocol transformation with our examples of protocol compositions.

Example 14 The transformed strands of the one-to-one protocol composition $NSL;_S DB$ of Example 10 are as shown below:

$$
\begin{aligned}
&:: r :: [nil| + (\textit{NSL.init}),\\
&\qquad + (pk(B, n(A,r); A)), -(pk(A, n(A,r); NB; B)), +(pk(B, NB)),\\
&\qquad - (\textit{DB.resp}(r\#)), +(\textit{NSL.init}(r\#) \;.\; A \;.\; B \;.\; n(A,r))]\&\\
&:: r :: [nil| + (\textit{NSL.resp}),\\
&\qquad - (pk(B, NA; A)), +(pk(A, NA; n(B,r); B)), -(pk(B, n(B,r))),\\
&\qquad - (\textit{DB.init}(r\#)), +(\textit{NSL.resp}(r\#) \;.\; A \;.\; B \;.\; NA)]\&\\
&:: r', r\# :: [nil| + (\textit{DB.init}(r\#)), -(\textit{NSL.resp}(r\#) \;.\; A \;.\; B \;.\; NA),\\
&\qquad + (n(B, r')), -(NA \oplus n(B, r'))]\&\\
&:: r\# :: [nil| + (\textit{DB.resp}(r\#)), -(\textit{NSL.init}(r\#) \;.\; A \;.\; B \;.\; NA),\\
&\qquad - (N), +(NA \oplus N), nil]
\end{aligned}
$$

The transformed strands of the protocol composition $NSL;_S\, KD$ of Example 11 are as shown below:

$$
\begin{aligned}
&:: r, r\# :: [nil| + (\textit{NSL.init}),\\
&\qquad + (pk(B, n(A,r); A)), -(pk(A, n(A,r); NB; B)), +(pk(B, NB)),\\
&\qquad + (\textit{NSL.init}(r\#) \;.\; A \;.\; B \;.\; h(n(A,r), NB))]\&\\
&:: r, r\# :: [nil| + (\textit{NSL.resp}),\\
&\qquad - (pk(B, NA; A)), +(pk(A, NA; n(B,r); B)), -(pk(B, n(B,r))),\\
&\qquad + (\textit{NSL.resp}(r\#) \;.\; B \;.\; A \;.\; h(NA, n(B,r)))]\&\\
&:: r' :: [nil| + (\textit{KD.init}), -(\textit{RO1} \;.\; C \;.\; D \;.\; K),\\
&\qquad + (e(K, skey(C, r'))), -(e(K, skey(C, r'); N)), +(e(K, N))]\&\\
&:: r' :: [nil| + (\textit{KD.resp}), -(\textit{RO2} \;.\; C \;.\; D \;.\; K),\\
&\qquad - (e(K, SK)), +(e(K, SK; n(C, r'))), -(e(K, n(C, r')))]\&
\end{aligned}
$$

where *RO1* and *RO2* are variables of sort Role. ■

8. Pragmatic and Experimental Evaluation

In this section we further explore composition via protocol transformation versus composition via synchronization messages comparing them for ease of use and simplicity. Furthermore, we present some experimental results about the performance of the two approaches. First, in Section 8.1 we show the attack for the NSL-DB explained in Section 2.1. Then we fix the NSL-DB protocol using a hash function, as explained in Section 2.1, and show that the protocol is verified as secure by our tool, i.e., the search space is finite and no attack is found. Moreover, in Section 8.2 we show that the NSL-KD protocol presented in Section 2.2 is also verified as secure by the Maude-NPA. Each time we show a protocol

secure, we also show that a regular execution can be performed, proving that the search space is not empty a priori; however, these regular execution proofs have not been included in this paper, though they are available online (see below).

Here, the reader can see that the attack state patterns associated to the transformed protocol are more complex and hence more error prone when they have to be specified than the attack state patterns for composition via synchronization messages, since the introduction of fresh variables for protocol composition has to be done manually. Also, the attack state patterns look more artificial in the protocol transformation because of the back and forth messages.

In Section 8.3 we provide more details of the experiments and compare the results obtained using both techniques. All the experiments, including the source Maude-NPA files and the generated outputs, can be found at: `http://www.dsic.upv.es/~sescobar/Maude-NPA/composition.html`

8.1. The NSL-DB Protocol

We start with the NSL-DB protocol composition. As explained in Section 2.1, this protocol has an attack in which the honest principal B thinks that he has heard from a principal D (who may or may not be honest), but who has actually heard from an honest principal A. This covers, for example, the case in which D is dishonest, and tries to pass on an honest principal's authenticated response as his own. This attack is represented in Maude-NPA by an attack state pattern, according to the protocol specification of Example 12, where: (i) the first strand is Alice talking to some principal C acting as NSL initiator and connecting to a DB responder, (ii) the second strand is Bob taking to some principal D acting as DB initiator and receiving data from NSL responder, and (iii) we include disequality constraints for principal names, namely $a \neq D$ and $C \neq b$.

More specifically, the attack state pattern using the protocol transformation technique is as follows:

$$
\begin{aligned}
&:: r :: [nil, + (\textit{NSL.init}),\\
&\qquad + (pk(C, n(a,r); a)), -(pk(a, n(a,r); NC; C)), +(pk(C, NC))\\
&\qquad \mid -(\textit{DB.resp}(r\#1)), +(\textit{NSL.init}(r\#1) \mathbin{.} a \mathbin{.} C \mathbin{.} n(a,r))] \;\&\qquad (\textit{NSL-DB-a0-PT})\\
&:: r', r\#2 :: [nil, + (\textit{DB.init}(r\#2)), -(\textit{NSL.resp}(r\#2) \mathbin{.} D \mathbin{.} b \mathbin{.} n(a,r)),\\
&\qquad + (n(b, r')), -(n(b, r') \oplus n(a,r)) \mid nil]\\
&\qquad \&((a \neq D), (C \neq b))
\end{aligned}
$$

And the backwards search from this attack pattern does not terminate[8] due to a state space explosion, and no initial state is found up to the depth reached by the analysis.

In protocol composition via synchronization messages the attack state pattern is as shown below:

[8] In [15] we reported termination, but this turned out to be a result of a bug in Maude-NPA's management of disequality constraints, which has since been corrected. The development of new semantics and implementation helped us to discover this bug.

$$
\begin{aligned}
&:: r :: [nil, +(pk(C, n(a,r); a)), -(pk(a, n(a,r); NC; C)), +(pk(C, NC)) \\
&\qquad \mid \{\textit{NSL.init} \rightarrow \textit{DB.resp} \,;;\, \text{1-1} \,;;\, (a; C; n(a,r))\}]\& \qquad (\textit{NSL-DB-a0-SM}) \\
&:: r' :: [nil, \{\textit{NSL.resp} \rightarrow \textit{DB.init} \,;;\, \text{1-1} \,;;\, (D; b; n(a,r))\}, \\
&\qquad + (n(b,r')), -(n(a,r) \oplus n(b,r')) \mid nil] \\
&\qquad \&(a \neq D, C \neq b)
\end{aligned}
$$

The backwards search from this attack state using composition via synchronization messages finds an initial state from which it is reachable, and thus demonstrates a distance hijacking attack. The exchange of messages of this attack is as explained in Section 2.1.

We then considered other attacks similar to the distance hijacking attack which however produced a smaller search space. In the following attack, we asked whether it is possible for an attacker to use an initiator A's nonce to participate in the distance-bounding part of the protocol without Alice having completed the corresponding NSL strand. The attack state is given below (note the different position of the vertical bars w.r.t. attack state *NSL-DB-a0-PT*);

$$
\begin{aligned}
&:: r :: [nil, +(\textit{NSL.init}), +(pk(C, n(a,r); a)) \\
&\qquad \mid -(pk(a, n(a,r); NC; C)), +(pk(C, NC)), \\
&\qquad\quad -(\textit{DB.resp}(r\#1)), +(\textit{NSL.init}(r\#1) \,.\, a \,.\, C \,.\, n(a,r))]\& \qquad (\textit{NSL-DB-a1-PT}) \\
&:: r', r\#2 :: [nil, +(\textit{DB.init}(r\#2)), -(\textit{NSL.resp}(r\#2) \,.\, D \,.\, b \,.\, n(a,r)), \\
&\qquad + (n(b,r')), -(n(b,r') \oplus n(a,r)) | nil] \\
&\qquad \&(a \neq D, C \neq b)
\end{aligned}
$$

This, besides being simpler, required only that the bar move one step forward in the NSL strand, and produced a smaller search space in which the protocol transformation version was able to find an attack, and to terminate on the corrected version of the protocol, giving us a better opportunity compare the performance of the two approaches. The same result is obtained for the attack pattern *NSL-DB-a1-SM* but we do not include it here.

As explained in Section 2.1, the distance hijacking attack can be avoided using a hash function. The previous property for the NSL-DB is specified in the new version of the protocol with the following attack state pattern using the protocol transformation:

$$
\begin{array}{l}
:: r :: [nil, +(\textit{NSL.init}), \\
\qquad +(pk(C, n(a,r); a)), -(pk(a, n(a,r); NC; C)), +(pk(C, NC)) \\
\qquad \mid -(\textit{DB.resp}(r\#1)), +(\textit{NSL.init}(r\#1) \,.\, a \,.\, C \,.\, n(a,r))] \,\& \qquad (\textit{NSL-DB-a0-fix-PT}) \\
:: r', r\#2 :: [nil, +(\textit{DB.init}(r\#2)), -(\textit{NSL.resp}(r\#2) \,.\, D \,.\, b \,.\, n(a,r)), \\
\qquad +(n(b,r')), -(n(b,r') \oplus h(D, n(a,r))) \mid nil] \\
\qquad \&((a \neq D), (C \neq b))
\end{array}
$$

However, as in the case NSL-DB protocol, the analysis using the protocol transformation does not terminate due to state space explosion and, thus, the security of the protocol for this attack state pattern cannot be proved.

The distance hijacking attack via synchronization messages is as follows:

$$
\begin{array}{l}
:: r :: [nil, +(pk(C, n(a,r); a)), -(pk(a, n(a,r); NC; C)), +(pk(C, NC)) \\
\qquad \mid \{\textit{NSL.init} \rightarrow \textit{DB.resp} \,;;\, \text{1-1} \,;;\, (a; C; n(a,r))\}] \& \qquad (\textit{NSL-DB-a0-fix-SM}) \\
:: r' :: [nil, \{\textit{NSL.resp} \rightarrow \textit{DB.init} \,;;\, \text{1-1} \,;;\, (D; b; n(a,r))\}, \\
\qquad +(n(b,r')), -(h(D, n(a,r)) \oplus (b,r')) \mid nil] \\
\qquad \&(a \neq D, C \neq b)
\end{array}
$$

The analysis of this protocol composition using the composition via synchronization messages, terminates finding no attack (see Table 1). Thus, the attack state is unreachable.

Therefore, we proceed in a similar way as we did before and provide an attack pattern with an earlier position of the vertical bar:

$$
\begin{array}{l}
:: r :: [nil, +(\textit{NSL.init}), +(pk(C, n(a,r); a)) \\
\qquad \mid -(pk(a, n(a,r); NC; C)), +(pk(C, NC)), \\
\qquad -(\textit{DB.resp}(r\#1)), +(\textit{NSL.init}(r\#1) \,.\, a \,.\, C \,.\, n(a,r))] \& \qquad (\textit{NSL-DB-a1-fix-PT}) \\
:: r', r\#2 :: [nil, +(\textit{DB.init}(r\#2)), -(\textit{NSL.resp}(r\#2) \,.\, D \,.\, b \,.\, n(a,r)), \\
\qquad +(n(b,r')), -(n(b,r') \oplus h(D, n(a,r)))|nil] \\
\qquad \&(a \neq D, C \neq b)
\end{array}
$$

The analysis of the protocol using the protocol transformation terminates, finding no initial state from which this more specific attack state pattern is reachable. The same result is obtained for the attack pattern *NSL-DB-a1-fix-SM* but we do not include it here.

8.2. The NSL-KD Protocol

For the NSL-KD protocol presented in Section 2.2 we may wish to guarantee that a dishonest principal is not able to learn the secret key of an honest principal. This property is represented by an attack state pattern, according to the protocol of Example 13, where the first strand is an initiator of the KD protocol generating the session key $skey(a, n(a, r'))$, the second strand is a responder of the KD protocol using the same session key $skey(a, n(a, r'))$, and we ask whether the intruder can learn this session key by adding the fact $skey(a, n(a, r'))$ to the intruder knowledge.

More specifically, in the protocol transformation the attack state pattern is of the following form:

$$
\begin{aligned}
&:: r' :: [nil, +(\textit{KD.init}), -(\textit{RO1} \;.\; a \;.\; b \;.\; K),\\
&\qquad + (e(K, skey(a, r'))), -(e(K, skey(a, r'); n(b, r))), +(e(K, n(b, r))) \mid nil] \&
\end{aligned}
\quad (\textit{NSL-KD-PT})
$$

$$
\begin{aligned}
&:: r :: [nil, +(\textit{KD.resp}), -(\textit{RO2} \;.\; b \;.\; a \;.\; K),\\
&\qquad - (e(K, skey(a, r'))), +(e(K, skey(a, r'); n(b, r))), -(e(K, n(b, r))) \mid nil]\\
&\qquad \&(skey(a, r') \in \mathcal{I})
\end{aligned}
$$

whereas for the composition via synchronization message is specified as follows:

$$
\begin{aligned}
&:: r' :: [nil, \{\textit{NSL.init NSL.resp} \rightarrow \textit{KD.init} \;;;\; \text{1-*} \;;;\; (a; b; K)\},\\
&\qquad + (e(K, skey(a, r'))), -(e(K, skey(a, r'); n(b, r))), +(e(K, n(b, r)))|nil] \&
\end{aligned}
\quad (\textit{NSL-KD-SM})
$$

$$
\begin{aligned}
&:: r :: [nil, \{\textit{NSL.init NSL.resp} \rightarrow \textit{KD.resp} \;;;\; \text{1-*} \;;;\; (b; a; K)\},\\
&\qquad - (e(K, skey(a, r'))), +(e(K, skey(a, r'); n(b, r))), -(e(K, n(b, r)))|nil]\\
&\qquad \&(skey(a, r') \in \mathcal{I})
\end{aligned}
$$

Here again the reader can see that the attack state pattern for the transformed protocol lacks some useful information about what is really happening, since we have two strands, each participating in different protocol composition, but no indication of what the possible compositions are. However, the attack state pattern for the composition via synchronization messages clearly shows that the two different one-to-many compositions that are possible for each strand.

In this case, the desired property is satisfied by the NSL-KD, since the analysis terminates using both the protocol transformation and the composition via synchronization messages techniques, finding no initial state for the attack state pattern described above.

8.3. Performance Comparison

In this section we show in detail the results of the experiments presented in Sections 8.1 and 8.2. Table 1 gathers the results of the analysis of these protocol compositions, i.e., (i) the composition of the NSL and DB protocols (NSL-DB), (ii) the composition of the NSL and the fixed version of the DB protocol (NSL-DB-fix), and (iii) the composition of the NSL and the KD protocols (NSL-KD). Note that for the NSL-DB and NSL-DB-fix protocols we consider the two attack state patterns shown above: the more generic, denoted as “a0”, e.g. NSL-DB-a0; and the more specific, denoted as “a1”, e.g.

	Protocol Transformation				Composition via SM				
Attack	Secure?	Finite?	Depth	States	Secure?	Finite?	Depth	States	SM / PT
NSL-DB-a0	?	No	10	3434	No	Yes	16	1337	-
NSL-DB-a1	No	No	16	1529	No	Yes	13	259	0.17*
NSL-DB-fix-a0	?	No	10	2650	Yes	Yes	19	1690	-
NSL-DB-fix-a1	Yes	Yes	17	273	Yes	Yes	16	103	0.38
NSL-KD	Yes	Yes	19	1486	Yes	Yes	16	652	0.44

Table 1
Experiments with sequential protocol compositions

NSL-DB-a1. For each protocol composition we provide the following information. For each technique, i.e., protocol transformation and composition via synchronization messages (referred as composition via SM in the table header), the column “Secure?” shows whether the technique sucessfully proved the protocol composition is secure, i.e. Maude-NPA generated a finite search space finding no attacks, or insecure, i.e, Maude-NPA found an attack. When Maude-NPA did not obtain a definite result, i.e., when the analysis did not terminate (e.g. because of an state space explosion) and no initial state was found up to the depth reached by the analysis, we write “?" in this column. The column “Finite?” indicates whether Maude-NPA generated a finite state search space or not, i.e. whether the analysis of such protocol composition terminated or not. The column “Depth” provides the depth of the analysis, i.e., the number of reachability steps performed by Maude-NPA until: (i) it generates a finite search space with no attacks in the case of a secure composition, (ii) it finds the attack in the case of an insecure composition, or (iii) the analysis finished before obtaining a definite result; whereas the column “States” shows the total number of states generated during the analysis up to the indicated depth. For the composition via synchronization messages, the column “SM / PT” shows the state space reduction as the number of states explored by the synchronization messages method (SM) divided by the number of states explored by the protocol transformation method (PT). When Maude-NPA did not obtain concluding results using the protocol transformation technique we write “-” in this column. In the case of the simpler attack for the NSL-DB-fix protocol (attack NSL-DB-fix-a1 in Table 1), marked with an *, we considered only the number of states generated until the first initial state was found with both techniques, since Maude-NPA could not generate a finite search space in the protocol transformation approach.

Regarding the execution time of the experiments, we note that we present these for the purpose of comparing the composition times rather than as the best possible times that can be achieved using our methods. The Maude programming language offers several levels of programming, including core Maude that provides the basic functionality of Maude, and the meta-level, in which Maude programmers can design new functionalities. Core Maude has been carefully optimized, and hence programs in core Maude run faster than programs at the meta-level. Our approach has been to first implement functionality at the meta-level, and then, when it is well understood, have it implemented in core Maude. Many of the features we use in this analysis, including narrowing modulo equational are still implemented in the meta level although work is ongoing in moving them to core Maude.

With this in mind, we present the execution times as follows. In the case of attack NSL-DB-a0, the analysis using the protocol composition failed to complete after several days, whereas using the composition via synchronization messages it completed in a little under 9 hours. For attack NSL-DB-a1 using protocol transformation, the tool did not complete, and took almost 2 days to find an attack, while it took 1/2 hour to find the attack when using synchronization messages. Attack NSL-DB-fix-a0 ran for several

days without finishing when using protocol transformations, while it completed after 6 days when using synchronization messages. For attack NSL-DB-fix-a1, the execution time was reduced from an hour and a half when using protocol transformations to 35 minutes when using composition via synchronization messages. For attack NSL-KD the tool took nine hours to complete using protocol transformations versus one and one-half hours using synchronization messages.

The reader may wonder why attack NSL-DB-fix-a0 took so much longer to complete than the other attacks, even for synchronization messages. Although we have not yet investigated the reasons in detail, we believe that it is because that attack makes the most extensive use of narrowing modulo exclusive-or, which is the most expensive operation.

In summary, protocol transformation fails to provide a definite result about the security of two of our experiments, namely the analysis of the NSL-DB and NSL-DB-fixed protocol compositions for the distance hijacking attack state pattern, whereas this problem does not occur with the composition via synchronization messages. Morevoer, composition via synchronization messages generates a finite state search space in all cases, whereas with protocol transformation this happens in only two cases. Moreover, in the case in which both complete or both find an attack, so that it is possible to compare performance directly, both state space size and time spent improved significantly for synchronization messages. SM / PT state space size ratios ranged from 0.17 to 0.44. Ratios for time spent were even more dramatic, ranging from 0.000868 to 0.389. Although we should be careful about drawing too many conclusions for such a small number of experiments, we believe that it is safe to conclude that composition via synchronization messages offers a significant improvement in both space and time efficiency.

9. Related Work, Lessons Learned, and Future Directions

9.1. Related Work

Our work addresses a somewhat different problem than most existing work on cryptographic protocol composition, which generally does not address model-checking. Indeed, to the best of our knowledge, most protocol analysis model-checking tools simply use concatenation of protocol specifications to express sequential composition. However, we believe that the problem we are addressing is an important one that tackles a widely acknowledged source of protocol complexity. For example, in the Internet Key Exchange Protocol [27] there are sixteen different one-to-many parent-child compositions of Phase One and Phase Two protocols. The ability to synthesize compositions automatically can greatly simplify the specification and analysis of protocols like these.

Now that we have a mechanism for synthesizing compositions, we are ready to revisit existing research on composing protocols and their properties and determine how we could best make use of it in our framework. There have been two approaches to this problem. One, called *nondestructive* composition in [10], is to concentrate on properties of protocols and conditions on them that guarantee that properties satisfied separately are not violated by the composition. This is often (although not always) applied to parallel composition. This is, for example, the approach taken by Gong and Syverson [21], Guttman and Thayer [26], Cortier and Delaune [7], Ciobâcă and Cortier [9], Groß and Mödersheim [22], and, in the computational model, Canetti's Universal Composability [4]. The conditions in this case are usually ones that can be verified syntactically, so Maude-NPA, or any other model checker would only be of use here to supply an experimental method for testing various hypotheses about syntactic conditions.

Of more interest to us is the research that addresses the compositionality of the protocol properties themselves, called *additive* composition in [10]. This addresses the development of logical systems and

tools such as CPL, PDL, and CPSA cited earlier in this paper, in which inference rules are provided for deriving complex properties of a protocol from simpler ones. Since these are pure logical systems, they necessarily start from very basic statements concerning, for example, what a principal can derive when it receives a message. But there is no reason why the properties of the component protocols could not be derived using model checking, and then composed using the logic. This would give us the benefits of both model checking (for finding errors and debugging), and logical derivations (for building complex systems out of simple components), allowing to switch between one and the other as needed. Indeed, we think that Maude-NPA is well positioned in this respect. For example, the notion of state in strand spaces that it uses is very similar to that used by PDL [6], and we have already developed a simple property language that allows us to translate the "shapes" produced by CPSA into Maude-NPA attack state patterns. The next step in our research will be to investigate this connection more closely from the point of view of compositionality.

9.2. Lessons Learned and Future Directions

Our work has also taught us much about the optimum strategies for extending Maude-NPA. First of all, although it is desirable to be conservative when extending the syntax and semantics, this should be done in such a way that the resulting semantics reflects the extended functionality in a natural way. Secondly, the use of message passing over the Dolev Yao channel is expensive computationally and can lead to state space explosion. Thus it should be used only when the properties of the Dolev Yao channel are actually needed. Thirdly, the use of an abstract semantics which is not actually implemented can be very helpful in assisting us to experiment with different implementation approaches in the tool itself. This allowed us to compare performance of different approaches while understanding their relationship to the abstract semantics. Thus we could be sure that we were not giving up correctness in order to obtain better performance, and we could understand to what degree we were losing expressiveness.

Finally, since the work of [18] we have discovered that sequential protocol composition is a key idea for several other applications in protocol specification such as protocol branching, secure communication channels, group protocols and protocols with global state memory. We believe that these applications can also be supported in Maude-NPA with extensions of the methods presented in this paper to support more expressive composition languages. We have performed a preliminary study of these applications but we leave for future work a deeper investigation on these topics.

References

[1] Martín Abadi and Cédric Fournet. Mobile values, new names, and secure communication. In Chris Hankin and Dave Schmidt, editors, *POPL*, pages 104–115. ACM, 2001.

[2] M. Anlauff, D. Pavlovic, R. Waldinger, and S. Westfold. Proving authentication properties in the protocol derivation assistant. In *Proc. of Joint Workshop on Foundations of Computer Security and Automated Reasoning for Security Protocol Analysis*, 2006.

[3] Mark Baugher, Thomas Hardjono, Hugh Harney, and Brian Weis. The group domain of interpretation. Technical report, RFC 3547, July, 2003.

[4] Ran Canetti, Yehuda Lindell, Rafail Ostrovsky, and Amit Sahai. Universally composable two-party and multi-party secure computation. In *STOC*, pages 494–503, 2002.

[5] S. Capkun and J. P. Hubaux. Secure positioning in wireless networks. *IEEE Journal on Selected Areas in Communication*, 24(2), February 2006.

[6] I. Cervesato, C. Meadows, and D. Pavlovic. An encapsulated authentication logic for reasoning about key establishment protocols. In *IEEE Computer Security Foundations Workshop, 2005*, 2005.

[7] Véronique Cortier and Stéphanie Delaune. Safely composing security protocols. *Formal Methods in System Design*, 34(1):1–36, 2009.
[8] Cas J. F. Cremers, Kasper Bonne Rasmussen, and Srdjan Capkun. Distance hijacking attacks on distance bounding protocols. In *19th Annual Network and Distributed System Security Symposium, NDSS 2012, San Diego, California, USA, February 5-8, 2012*. The Internet Society, 2012.
[9] Ştefan Ciobâcă and Véronique Cortier. Protocol composition for arbitrary primitives. In *CSF*, pages 322–336. IEEE Computer Society, 2010.
[10] A. Datta, A. Derek, J. C. Mitchell, and D. Pavlovic. Secure protocol composition. In *Proc. Mathematical Foundations of Programming Semantics*, volume 83 of *Electronic Notes in Theoretical Computer Science*, 2003.
[11] Y. Desmedt. Major security problems with the "unforgeable" (Feige-)Fiat-Shamir proofs of identity and how to overcome them. In *Securicom 88, 6th worldwide congress on computer and communications security and protection*, pages 147–159, Paris France, March 1988.
[12] S. Doghim, J. Guttman, and F. J. Thayer. Searching for Shapes in Cryptographic Protocols. In *TACAS 2007*. Springer LNCS 4424, March 2007.
[13] D. Dolev and A. Yao. On the security of public key protocols. *IEEE Transaction on Information Theory*, 29(2):198–208, 1983.
[14] Francisco Durán and José Meseguer. A maude coherence checker tool for conditional order-sorted rewrite theories. In Peter Csaba Ölveczky, editor, *WRLA*, volume 6381 of *Lecture Notes in Computer Science*, pages 86–103. Springer, 2010.
[15] N. Durgin, J. Mitchell, and D. Pavlovic. A Compositional Logic for Program Correctness. In *Fifteenth Computer Security Foundations Workshop — CSFW-14*, Cape Breton, NS, Canada, 11–13 June 2001. IEEE Computer Society Press.
[16] S. Escobar, C. Meadows, and J. Meseguer. A rewriting-based inference system for the NRL protocol analyzer and its meta-logical properties. *Theor. Comput. Sci.*, 367(1-2):162–202, 2006.
[17] S. Escobar, C. Meadows, and J. Meseguer. Maude-NPA: Cryptographic protocol analysis modulo equational properties. In A. Aldini, G. Barthe, and R. Gorrieri, editors, *FOSAD 2008/2009 Tutorial Lectures*, volume 5705 of *LNCS*, pages 1–50. Springer, 2009.
[18] Santiago Escobar, Catherine Meadows, José Meseguer, and Sonia Santiago. Sequential protocol composition in Maude-NPA. In Dimitris Gritzalis, Bart Preneel, and Marianthi Theoharidou, editors, *ESORICS*, volume 6345 of *Lecture Notes in Computer Science*, pages 303–318. Springer, 2010.
[19] Santiago Escobar, Ralf Sasse, and José Meseguer. Folding variant narrowing and optimal variant termination. *J. Log. Algebr. Program.*, 81(7-8):898–928, 2012.
[20] F. J. Thayer Fabrega, J. Herzog, and J. Guttman. Strand Spaces: What Makes a Security Protocol Correct? *Journal of Computer Security*, 7:191–230, 1999.
[21] L. Gong and P. Syverson. Fail-stop protocols: An approach to designing secure protocols. In R. K. Iyer, M. Morganti, W. K. Fuchs, and V. Gligor, editors, *Proc. of the 5th IFIP International Working Conference on Dependable Computing for Critical Applications (Urbana-Champaign, IL, Sept. 1995)*, pages 79–99. IEEE Computer Society Press, Los Alamitos, CA, 1998.
[22] Thomas Groß and Sebastian Mödersheim. Vertical protocol composition. In *CSF*, pages 235–250. IEEE Computer Society, 2011.
[23] J. Guttman. Security protocol design via authentication tests. In *Proc. Computer Security Foundations Workshop*. IEEE Computer Society Press, 2001.
[24] J. D. Guttman, J. C. Herzog, V. Swarup, and F. J. Thayer. Strand spaces: From key exchange to secure location. In Carolyn Talcott, editor, *Workshop on Event-Based Semantics*, 2008. Position papers available at `http://blackforest.stanford.edu/eventsemantics/`.
[25] Joshua D Guttman. Security goals: Packet trajectories and strand spaces. In *Foundations of Security Analysis and Design*, pages 197–261. Springer, 2001.
[26] Joshua D. Guttman and F. Javier Thayer. Protocol independence through disjoint encryption. In *CSFW*, pages 24–34. IEEE Computer Society, 2000.
[27] D. Harkins and D. Carrel. The Internet Key Exchange (IKE), November 1998. IETF RFC 2409.
[28] Catherine Meadows. Analysis of the internet key exchange protocol using the nrl protocol analyzer. In *Security and Privacy, 1999. Proceedings of the 1999 IEEE Symposium on*, pages 216–231. IEEE, 1999.
[29] Catherine Meadows and Dusko Pavlovic. Deriving, attacking and defending the gdoi protocol. In *Computer Security–ESORICS 2004*, pages 53–72. Springer, 2004.
[30] Catherine Meadows, Paul Syverson, and Iliano Cervesato. Formal specification and analysis of the group domain of interpretation protocol using npatrl and the nrl protocol analyzer. *Journal of Computer Security*, 12(6):893–931, 2004.
[31] J. Meseguer. Conditional rewriting logic as a unified model of concurrency. *Theoretical Computer Science*, 96(1):73–155, 1992.
[32] J. Meseguer. Membership algebra as a logical framework for equational specification. In F. Parisi-Presicce, editor, *Proc. WADT'97*, pages 18–61. Springer LNCS 1376, 1998.

[33] Michael Rusinowitch and Mathieu Turuani. Protocol insecurity with a finite number of sessions and composed keys is NP-complete. In 14^{th} *IEEE Computer Security Foundations Workshop*, pages 174–190, 2001.
[34] J. Meseguer S. Escobar, C. Meadows and S. Santiago. Sequential protocol composition in Maude-NPA. Technical Report DSIC-II/06/10, Universitat Politècnica de Valéncia, June 2010.
[35] Sonia Santiago. *Advanced Features in Protocol Verification: Theory, Properties, and Efficiency in Maude-NPA*. PhD thesis, Universitat Politècnica de València, January 2015.
[36] TeReSe, editor. *Term Rewriting Systems*. Cambridge University Press, Cambridge, 2003.
[37] P. Thati and J. Meseguer. Symbolic reachability analysis using narrowing and its application verification of cryptographic protocols. *J. Higher-Order and Symbolic Computation*, 20(1–2):123–160, 2007.
[38] Patrick Viry. Equational rules for rewriting logic. *Theor. Comput. Sci.*, 285(2):487–517, 2002.